1
Divide and conquer

1.1	Example 1: CDROM design	4
1.2	Theory 1: Multiple estimates	7
1.3	Theory 2: Tree representations	10
1.4	Example 2: Oil imports	13
1.5	Example 4: The UNIX philosophy	15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought...endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2,789,164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

Is there a reason other than simplicity for estimating the CD's area as a square rather than a circle? If we overestimated the length of a side (or accurately estimated it) we would have a bigger error.

> I'm guessing that it's just so you can do the calculations easily without a calculator.
>
> I'm fairly sure it has to do with using pi. If you took the estimate of 10, and did 5^2 * 3 (for pi), you get 75. 10^2 is 100. The actual is 6^2 *3.14 which is 113. Obviously if it were smaller, it would miss by a larger amount, but you can probably deduce that for yourself and make that judgement.
>
> > It just feels so counterintuitive to just square the lengths of the cd when we've been taught that the area is pi*r^2. Can't we just multiply by 3 instead of not multiplying by pi at all?
> >
> > > I think that all we're doing here is assuming that the area of a circle is about the same as the area of a square with a side length slightly less than the circle's diameter, which makes our lives easy and doesn't require multiplication by anything but 10s.
>
> I think simplicity is the sole motivator, which is especially justifiable in estimation – no?

Is it alright that in the estimations I would have picked quite different numbers? For instance, in estimating sampling rate, I have read that most adults cannot hear frequencies over 16kHZ so I probably would have estimated 15 instead of 20 kHZ. Is it more important to have the right numbers or is the emphasis more on being able to explain your reasoning?

> Given the inherently imperfect nature of estimation, I would imagine logic is more important than actual numbers. Of course, the answer must make some sort of sense to be useful.

Are we supposed to answer the questions on page 6 (problems 1.1 and 1.2)? If so, I do not understand what problem 1.1 is asking, are there any definitions I am missing or am I just not reading it right?

> You don't need to formally answer those inline problems (sorry for splitting the infinitive). They are placed there to help you review ideas in the preceding material. I'll add more of them as I revise the textbook. I hope that these questions help people use the textbook for self study without an "Art of Approximation" course at their university (or they may be long graduated and, say, practicing as an engineer).

1
Divide and conquer

1.1	Example 1: CDROM design	4
1.2	Theory 1: Multiple estimates	7
1.3	Theory 2: Tree representations	10
1.4	Example 2: Oil imports	13
1.5	Example 4: The UNIX philosophy	15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have $2,789,164$ files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

The suggestion on pages 6-7 to break the numbers down into powers of ten and others was extremely helpful to me!! I think having this thought in the back of my head for this class (and for life) will prove to be very useful!!

This link didn't work on my apple laptop

It is working on mine...

NB requires Firefox. Is that the cause?

I don't even think you can comment without Firefox (based on my try with Safari), so they probably have a different issue.

I found it had problems when I was on a slower WiFi network. The scripts ran too slowly and the whole site would fail to run. Once on a landline, everything was fine.

I'm also curious as to whether it is more important to start with "correct" approximations or be able to explain your reasoning throughout an estimation. I was confused about sampling rate and some terminology as well when I first read this last night but after reading through the notes most of my questions have been answered.

For the purposes of the class, it seems the explanation is much more important than the immediate result. It is hard to ignore the fact, however, that the initial approximations do have a pretty profound impact on the result. So perhaps the skill comes from minimizing that impact through approximation methods?

I am curious about this also. It seems in a lot of your examples in class, high guesses and low guesses cancel out, but even for that to work you have to be at least reasonable. Is this just from practicing guessing so much? Reasoning for the numbers is obviously important, but what I would have come up with for these examples would have been way off - too far off for my reasoning to recover my guess.

For question 1.1 I think a general factor is something globally applicable, like inches in a foot or seconds in an hour. A particular form of unity is specific to one case. Here the samples per second we calculate is based on what we know about the CD. In this case, sample rate and playtime, and sample size. If any of these were changed we would have a different conversion.

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have $2,789,164$ files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

Responding to a comment above, I think that w/ the computer the idea of divide and conquer is valid because we divide up our population of files into manageable pieces the same way that civilizations have divided up people the same way in order to successfully manage them.

I think a sanity check for the CD example is that CD ROMs use IR/Red lasers to read the discs. If we assume that the spacing of the pits is on the order of a wavelength, then we know that the spacing should be around 1 um

I don't think I would have recognized to divide this problem into three parts (playing time, sampling rate, sample size) as outlined here, and I'm not sure my estimates for these values would have been at all accurate. How accurate are our solutions supposed to be (ie, within an order of magnitude?)

Yes, would a problem on the test have told us to divide it into these three different subsections as a hint or would we have had to figure it out on our own?

Interesting that the numbers generated here match the first solution almost exactly; I agree with previous comments that the first method seems much more intuitive and simpler. I feel like the second method is almost geared towards having a matching solution.

I also like the first method much better. The sampling rate and so forth was much more difficult for me to understand.

I agree with above, the sampling rate requires a great deal of outside knowledge which complicates the example more than is necessary. I would have to say the same about sample size.

is it really necessary to evaluate all of this through mental math?

I agree, this seems to be a very thoughtful estimation, why should we continue to solve this problem through mental math rather than just putting pencil to paper? Don't we want to preserve our estimations by giving an exact approximation? I know we get roughly the same answer wither way, but the straightforward math seems to eliminate a step of uncertainty in order to prevent "preventable" math errors?

It seems as if the assumptions being made in order to make the overall estimate. Is there a way to make an accurate estimate not knowing some of these assumptions?

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

I don't understand how this helps answer the question.

where is this comment pinned?

There seems to be a lot of assumptions that the reader has previous knowledge. A bit more explanation on where the numbers come from and why certain approximations are valid would be helpful. I also didn't understand why there was a conversion at the end. I thought we were looking for space between pits not bits/s.

I agree with this - a lot of these numbers assume prior knowledge (Beethoven's Ninth Symphony, 44kHz, sample requires 32 bits) which many people do not have - where do we turn to if we don't have/know this information?

Is it ever okay to just look up a figure if we really have no idea?

yeah I'm also unclear as to what a pit is.

A pit refers to the physical deformity of the cd on a microscopic level. The data has to be encoded in binary. The 1's and 0's correspond to pits and no-pits.

What is an antialias filter?

http://en.wikipedia.org/wiki/Anti-aliasing

Here you use the diameter of the CD to calculate the spacing between the pits, I don't really understand why you pick the diameter, because the distance between two points in a circle might be smaller than the diameter, so is this method going to overestimate the spacing?

It seems like the propagation of error would be a problem. Is is better to be more accurate in the first place or to understand how to compensate for error? It would be interesting to learn how to identify which aspects lead to the largest errors and which can be rougher estimates.

1
Divide and conquer

1.1	Example 1: CDROM design	4
1.2	Theory 1: Multiple estimates	7
1.3	Theory 2: Tree representations	10
1.4	Example 2: Oil imports	13
1.5	Example 4: The UNIX philosophy	15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

Read the intro and the first section (i.e. section 1.1). The reading memo is due (via NB) by 10pm on Thursday Feb 4. Have fun!

how are things "due" - should we just make comments?

> Right, just make comments.
>
> Then at 10pm on Thursday evening, I'll do two things: (1) Look at the reading page by page, going through all the comments to help me prepare the lecture; (2) Look at the comments student by student so that I can give everyone the 'decent effort' mark (I have to write a few programs to help me with that, but they are almost done).
>
> But feel free to comment after 10pm on Thursday as well.
>
> On the feedback sheets, several people asked, "What if our comments have already been made by others? How do we get credit for our effort?" The most important is that you can contribute by responding to other students' comments on other areas where you feel able to make suggestions.
>
> Discussing is the best way to learn.

Do we need to "sign" with our name or can we stay anonymous?

> You are free to stay anonymous. As the instructor, I can still know the author (except that I cannot see comments that you mark as just for yourself [the "Myself" radio button]).
>
> The NB default for a comment is that it will be visible to the class but that it is anonymous. I think that's a sensible default: It allows the whole class to benefit from the question, but doesn't force you to "expose" yourself if you do not want to (same principle as the paper feedback sheets at the end of lecture).
>
> As you no doubt have noticed, you can change the setting for any particular comment by selecting the appropriate radio button for the visibility ("class", "staff", or "myself") and can sign it if you choose by checking the "Sign" box.

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

Is there a way to separate only my notes and their replies?

What should be done if there isn't much to common on, like if it's fairly easy to understand the piece and it seems logical? Should we just read back and common on other comments?

Also, what if you comment a bit late, will that affect the grade? (I'll probably ask this in class too). For example, I read this last night and made two comments. Then, I decided I wanted to add more comments and added 1 at 9:55, 1 at 10pm exactly, and 2 at 10:03. Will having two at 10:03 make the grade a 'late' one, or will those comments just be counted as 'extra' comments, like additional ones just to help the paper. Thanks!

1

Divide and conquer

1.1	Example 1: CDROM design	4
1.2	Theory 1: Multiple estimates	7
1.3	Theory 2: Tree representations	10
1.4	Example 2: Oil imports	13
1.5	Example 4: The UNIX philosophy	15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

I REALLY don't like this software. Is this our only option for this class?

> I'm sorry to learn that. But please would you explain what aspects you do not like? The software, which is written by Sacha Zyto in CSAIL, is being improved all the time. Many of the improvements are based on suggestions from users (often from dissatisfied users). So, your particular suggestions, complaints, and comments will be very welcome.

Is there a way to see all the comments at once? Some of the older ones are not visible at all, and to see more of them you have to scroll each page (which is fine, but if you have already read it, it would be cool to just be able to scroll through the comments, and when you want to see where it was talking about, click to have it take you there, just like it works now). Sorry if I missed how to do that in the tutorial.

> Sacha (the author of NB) is working night and day on the revised user interface, and the view that you request will be part of it. Till then, he offered a temporary fix, which is this URL: <http://nb.csail.mit.edu/?t=p20&view=auth_gro
>
> I will place a link to it from the 6.055 course website <mit.edu/6.055/> so that you do not have to copy and paste that URL.
>
> The comments are displayed in the order of most popular (number of people who clicked on the box saying "This is unclear" or "Answer, please!"). I'd like to see it ordered by pages, but maybe that will come with the new user interface.

In the "How To Do Reading Memos" article, it explains that the notes are meant to aid in our understanding, while helping the author. With 75 people commenting on the same 5 pages, it's almost impossible to follow any one persons (including your own) thought process. Also, the comments are very difficult to read, especially when they are not in order based on what they are referring to in the text. When you click on some comments the reference is not highlighted. The "entire article" notes appear each time you click on a new page so you end up reading the same material over and over again. I wasn't able to post a comment on the whole article as others have. I wasn't able to post the "clear now" or answered" status on comments. Each time I try to edit a comment, the menu shows up and blocks the entire text entry field. I would prefer to write comments on a paper copy and ask questions from others in a forum.

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought... endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

Thank you! I can answer perhaps one or two of your questions. To comment on the whole article, click on the middle icon on the top left (in the "Nav" tab). If you hover over it, it will say, "Add a comment on the whole file", and the icon itself looks like a slanted pencil.

About the comments not being in the order based on the text: I'm working on a system for myself that will make a PDF file with comments on the left half of the page and the original PDF page on the right half – so that I can prepare for lecture by reading through the text and comments in order. If that works, I'll make it available on the website so that anyone in the class can get the same information.

For the other questions, I'm going to defer to higher authority and ask Sacha to respond himself...

I understand that that's the button to use, but the site doesn't respond when I click on it. Any other improvements would be a welcome relief.

well, nobody is using the "answer please" or "unclear" buttons, so they are essentially useless, and there is no organization to where the comments lie.

sorry if any of these comments have already been addressed, i'm finding it hard to read through all the comments, so i appoligize if i duplicate others comments.

1. i'd like the ability change the FONT size of the txt in the comments section so i can see more words at one time on my screen. 2. i'd like to have "pop out feature" so i can read just the pdf without the distraction of the other stuff on the pg 3. i'd like the pictures on the RT section to scale in size when i resize the section, as it is i loose the ability to zoom in if i make the rt section small. 4. when i resize this comment section i loose my comments, ie if i try to resize it before i click the save button, i ahve to retype the comment.

ty for being open to comments, delighted to hear someone is working on improving the sw.

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

Hello everyone, this is Sacha and I'm one of the NB designers. Thanks a lot for your feedback ! This class has been participating so well, that we didn't initially plan to have so many notes on a page. Hence, we've been working on a new UI, which we'll hopefully release next week. In the meantime, if you want to see an **unfinished untested alpha preview** (i.e. if it can make your reading over the week-end reading easier), check out: http://nb.csail.mit.edu:8080/?t=p23 (note different port and URL parameter).

It's a bit confusing to understand. Does everyone show up as Anonymous?

Only if he/she chooses. The default choices are that the note is visible to the class but that it is unsigned (anonymous). To sign your note, check the "Sign" box before you save the note.

But as the instructor I can see the author name, so you can freely make your notes anonymous and get proper "decent effort" credit.

I agree with some of the issues raised here - it's difficult to navigate so many comments at once, and I can't make a comment without a toolbox hovering over my textbox so that I cannot see what I'm typing. It would be nice to have a n option to enlarge the comment area or perhaps show only comments that are relevant to the part of the paper that you are viewing. Additionally, if you scroll through the pages manually using the sidebar on the right side of the screen, the page number in the upper left hand corner does not update (the part where it says 1/5, etc). This causes it to spaz sometimes when you write a new comment. All that being said, the program is still quite impressive, and I'm sure that with some tweaking it could become a very useful class tool.

Good to know. I hadn't noticed that option until now!

Is it really "history" what is helping us solve problems of our time? A better way to write this sentence is: "Throughout history, a common strategy (referring to "divide-and-conquer") has been successful and still is today." The strategy has proven to be good throughout history, but it is not "history" itself what is helping us solve problems of our time.

What is meant exactly by dominion? It seems as if any subject under another's dominion would be resentful

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

is brute force really that costly? after all, soldiers in history are generally poorly paid. I understand the idea being given here, but the wording lacks precision.

I love this. Most of the time text books start out with some 'practical' story, but it's always dry and useless. I get, now, where the idiom 'divide and conquer' comes from. Thank you!

Does this statement represent the idea that resentful subjects would expend their energy fighting one another? It seems as if those in power are causing the fighting to take place.

I had never seen anyone clarify this sort of thing in this manner. It would be more appropriate to say "in prose" in a complete sentence: "I have stressed the relevant words in italics..."

> no, that's standard practice for block quotes

I always thought that the divide and conquer was speaking to the enemies not the subjects in a kingdom, but it makes sense when you think about it

I wonder: is it possible to mathematically model the optimal size of an empire, given a unit tribe size and length of time said tribe has been with the empire?

> I think it would be very difficult to model this especially going into the future. Due to modern communications, I think it is getting difficult to divide and isolate different "tribes"
>
> > I'm a bit curious to know how ancient empires 'modeled' this effectively as well. Tribes in real life can't be divided so evenly, and would likely be more along cultural lines. Perhaps they DIDN'T do so well - that may explain why some civilizations do better than others? Can approximation really be used for something as complex as this?

I would think this is more a strength. Don't most terrorist tactics revolved around organizing into separate & distinct cells?

Convincing tribes to fight one another made me think of recursions in programming, which i see now is just divide and conquering a calculation. Which I guess is also what we're attempting to do in approximation. Though answers to smaller solutions don't automatically create more solutions, they can allow us to realize things that weren't at first apparent.

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought...endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

While solving the problem of discomfort, this is probably going to create the problem of a "less mighty" empire which interferes with an empire's original goal (usually)

What do we do if we don't have any questions or comments? I generally understand the reading and don't have much to say that would benefit other people.

> You can be very helpful by answering questions that other students have raised (and there are so many questions to choose from, far too many for me to answer by myself!).

It's interesting that divide an conquer is a tactic both used in army fighting, interview questions, and many many important algorithms in computer science, like matrix multiplication for example.

Perhaps it would be more effective to give a range for the number of files a typical person has on his/her computer?

1
Divide and conquer

1.1 Example 1: CDROM design 4
1.2 Theory 1: Multiple estimates 7
1.3 Theory 2: Tree representations 10
1.4 Example 2: Oil imports 13
1.5 Example 4: The UNIX philosophy 15

How can ancient Sumerian history help us solve problems of our time?

From Sumerian times, and maybe before, every empire solved a hard problem – how to maintain dominion over resentful subjects. The obvious solution, brute force, costs too much: If you spend the riches of the empire just to retain it, why have an empire? But what if the resentful subjects would expend their energy fighting one another instead of uniting against their rulers? This strategy was summarized by Machiavelli [19, Book VI]:

> A Captain ought. . . endeavor with every art to divide the forces of the enemy, either by making him suspicious of his men in whom he trusted, or *by giving him cause that he has to separate his forces*, and, because of this, become weaker. [my italics]

Or, in imperial application, divide the resentful subjects into tiny tribes, each too small to discomfort the empire. (For extra credit, reduce the discomfort by convincing the tribes to fight one another.)

Divide and conquer! As an everyday illustration of its importance, imagine taking all the files on your computer – mine claims to have 2, 789, 164 files – and moving them all into one directory or folder. How would you ever find what you need? The only hope for managing so much complexity is to place the millions of files in a hierarchy. In general,

I think there's a difference between just "divide and conquer" and the hierarchy of computer folders. The hierarchy itself contains information, which is lost when everything is tossed into the same folder.

I disagree that this is the only hope. Search methods like Spotlight or Google avoid the divisions of hierarchy, which can actually obscure connections between data. As pointed out, hierarchy has other advantages for binning, but not necessarily for finding.

However, it is much easier to find things (if you don't want to use a search engine, or if you don't have the exact filename) if you do use the divide and conquer concept. For things like image files, unless you know what the file name is, it makes much more sense to look for it with a certain timestamp, which is a way to divide up the images. Sure it's not the "only hope," but its a good way to generalize

I also think this is not the "only hope" but I do think there are very distinct and key advantages that dividing and conquering provides that spotlight or google cannot (for example if you do not remember the file name but remember it was for your 2.038 class....having a 2.038 folder would come in handy)

I wonder what the memory usage and time usage are for each different type of search/filtering?

I agree that it is not quite the same to "divide and conquer" as to organize something, unless you see it from a very broad abstract level, but it might be too early in the chapter to make a connection like this. There is a stupid context menu that won't let me see what I type!!!!!!! and it wont go away!.

From what I have learned about divide and conquer, I don't see how its a hierachery. I assume every component you use is of importance you just have to figure out how to relate them to each other. Or do you mean the triangular way of dividing for divide and conquer is like a hierarchy.

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

How far apart are the pits on a compact disc (CD) or CDROM?

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

some societal pieces that emanated from imperial execution of divide and conquer are not all that manageable. I am from a country where this is a salient concern due to the fragmented nature of the the communities.

It seems like a fine analogy to me.

Perhaps the computer folder analogy isn't the best. It seems that rather than dealing with everything at once, dive and conquer breaks things apart as it goes, simplifying and approximating for ease. It's not divide everything, just whatever's at hand.

I agree..I don't think the problem of finding a file is made easier just by separating the files. It's important to have some classification of the groups when you order them so you can discard those groups that aren't of interest to you.

might sound better if you move "problems" to the end of the sentence, after "scientific".

I really like this summary you did in this memo..I have noticed it in later memos lately. I think its easier for me to understand the relation between the memo and example because sometimes I am completely lost.

Since you mentioned the examples, it might be useful to mention or introduce the 2 discussions here as well.

I have a few comments about this CD example in general. To be honest, I can't even remember the last time I handled or used a CD. I wouldn't have been able to estimate the information given regarding the physical size of the CD and its storage capacity. I would have had to look up all that information, which sort of ruins the estimation experience. And while I might have been able to guess at the physical size, I certainly would have no idea when it came to information like sampling rate or sample size. So while I think it was a good example, I think a lot of what was trying to be communicated here was lost because so many of these calculations seem to hinge on the ability to recall information about a CD.

This assumes prior knowledge of how a CD stores info, which is not necessarily common knowledge

Maybe you can start with a simpler example in stead of jumping right into an EE example.

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

How far apart are the pits on a compact disc (CD) or CDROM?

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

I'm unfamiliar with the terminology here. I'm guessing pits are the grooves cut into a cd, but I feel that should be more explicit.

> Thanks and agreed. Each bit (0 or 1) is stored as a tiny hole on the CD. I think the "0" hole has a slightly different shape than the "1" hole. The pits are indeed arranged in a spiral, sort of like grooves on an old-style vinyl record. So there are really two distances involved: The distance between pits (by which I really mean between the centers of the pits) along the spiral, and the distance between pits on neighboring tracks of the spiral. To keep the estimates simple, and doable mentally, I ignore all of the above and just imagine that the pits make a square lattice and that the CD itself is a square.
>
> A wise man said: "It is better to be approximately right than exactly wrong." (John Tukey)
>
> Another wise man said: "The art of being wise is knowing what to overlook."
>
> Those quotes are one of the main themes of the whole course.

Once you've said to approximate as a square lattice, this makes a lot of sense to me. But until you "give me permission to do so" I would never know that it's okay. This is something I hope to learn in this class.

Actually, there aren't '1' pits and '0' pits; rather, a change from pit to no pit indicates a 1 and no change indicates a zero. I bring this up because you can actually have up to 10 consecutive zeros in a row, which would be 10 bits of data on one pit or one non-pitted area, so our estimate of the number of pits could be significantly off.

See here: http://en.wikipedia.org/wiki/CD#Physical_details

So is there a set distance for either a pit or non-pit to occur? How does the reading device know how far to go before assuming that the there is no pit and therefore should input a zero

Agreed - I think pictures/diagrams would be extremely useful here.

My first thought was, "Why do we want to find out how far apart the pits are? Why is it important for us to know this information." For a first example, perhaps it would be more useful to provide one that the reader can actually relate to?

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

How far apart are the pits on a compact disc (CD) or CDROM?

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

I don't really know what a pit is.

It's a tiny hole made by scooping out a tiny bit of the CD material (usually a plastic). Depending on the shape of the hole, or pit, the CD player's laser will read it as a "0" or a "1".

Is this pit on the outer-most surface of the CD? if so, how is it protected from abrasion, and if it is not, how does it get scooped out? Especially with re-recordable disks.

I had no idea that CD's worked like this. I assume there's a layer of coating that exists to prevent damage to these pits.

How do rewriteable discs work? The little holes would have to be filled out and replaced, right? Also, how do dual layer discs work?

How would one consider dual-layer discs, or Blu-rays for example? Merely just another step of dividing and conquering can solve this?

The section just began, and the problem is already divided into pieces. I lose the big picture here. Although the following derivation is easy to follow, I sometimes forget why we are doing what we are doing. It might be good to outline how we can break up the problem and the different alternatives for solving each part. A diagram might be very useful to show the breakup of the problem.

Could you clarify how pits relate to a CD's data capacity?

From inference, it appears that pits on CDs store data.

Without understanding how pits are related to data storage this is a somewhat confusing leap. It would probably help to clarify how pits work at the beginning so that students don't have to backtrack and reevaluate after it's learned that each pit encodes one bit.

As near as I can tell, the information is stored as a binary code, 0's and 1's. So each pit and the sequence of pits encodes a tiny bit of data, then when read out in the proper order yields up whatever was stored to the CD

Each side? I thought we were talking about CD's...

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

How far apart are the pits on a compact disc (CD) or CDROM?

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

You are introducing a new and unexpected concept here: approximation. It makes no sense at all to approximate the area of a circle as the area of a square unless you don't have a calculator at hand. In fact, if you assume that the diameter is 10 cm, getting the area is simply 314/4 which equals roughly 75. That would be a better example of estimation. Before making approximations, you must make an argument for approximations in the text. If you dedicate a prior chapter to this, then you should remind us that we are trying to learn how to approximate. It actually took more text to justify an unneeded approximation than deriving a more exact answer with A=pi*(r1^2-r0^2).

It seems like the approximations made using a square to represent the circular cd are a bit inaccurate. How do we decide how much accuracy is worth a certain amount of time? I think I could find the circular area with more precision in a bit more time.

How can you know that (10cm)^2 is a good approximation of the area of the CD without actually measuring the area of the hole in it? It seems to me that making this approximation comes from the subtraction of the hole area from the CD area. But if we measure these areas, doesn't it stop being an approximation?

It only becomes clear later that the reason for estimating as a rectangle is to estimate the CD as a lattice of pits (rather than a spiral). When presented here, it's a bit: "uh, okay, but pi*r^2 is just as easy".

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

How far apart are the pits on a compact disc (CD) or CDROM?

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

does it really make it that much simpler to estimate the cd's area as a square? or is there some other reason

I mean, the numbers come out to within the same order of magnitude (100 vs 3*5^2 or 3*6^2) and in this case because of the large whole in the middle, using a square is even easier (compared with 3*(6^2-.75^2)). Plus, it makes the analogy of overlaying a lattice of pits to calculate the spacing between pits easier to understand/appropriate.

Would it be more "accurate" (quotes since this is an approximation anyway) to assume everything is in a square lattice on a square and then scale to an approximate area of the circular disk? The square lattice method of picturing the pit spacing is still valid, but the area approximation is a little closer to actual, and not really harder to calculate.

I'm still unclear about what kinds of generalizations we are allowed to make under which circumstances... at what point do we become more concerned with the order of magnitude than with a more accurate approximation?

I agree. I guess we're all too obsessed with exactness because we've been trained to think like this, but what's the standard criterion for level of accuracy? At what point is an approximation no longer really acceptable?

It does seem like the approximation is a little vague - we know the diameter of the CD is 12cm, but after that it's like we just threw off 2cm to make our numbers nice - at what point is that "acceptable" approximating, and at what point is that just guessing?

The example kind of justifies rounding down to ten, saying it accommodates for the hole in the middle of the disk that has no pits.

I wonder about whether or not (& how) the circular read/write constraint affects the usable area on the disk... Also, it seems important to note that capacity is broken down to the lowest discrete level (bit) that makes sense. This seems to allow the most fidelity in the following analysis.

I take back my comment from before.. this logic makes perfect sense to me.

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

▶ *How far apart are the pits on a compact disc (CD) or CDROM?*

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

is this something we would be expected to know off the top of our heads?

I agree. For whatever reason, the actual calculation seems much more straightforward and intuitive to me than the estimation. Perhaps it is because I don't know about sampling rate or sampling size but I don't feel like problem is one I would have been able to break down on my own.

Additionally, this doesn't take into account the space reserved for checking the data (making sure it isn't read incorrectly). This is why audio disks are hard to read when scratched - they lack the ability to check the data, so if part is scratched you're out of luck.

Shouldn't bits used for error checking still come out of the 700 megabytes? Isn't that just a function of the encoding method?

So, as the line says, he found this number by looking at a box of CDs that he had lying around...my guess is that he doesn't know this off the top of his head either. <3

While 700 mb is pretty stardard, a lot of cheaper CDs are 650... this is a number I have memorized from past experiences, so if i were calculating this, I probably would have used a number lower than 700.

So something like this perfectly fine to look on a box for the estimating. Is there ever a time when looking at a box is not reasonable or cheating? Can you ever really cheat in approximating?

I am not very techy, and this is something I would not have known.

I wouldn't have been able to do this, even with good estimation skills, because my nerd-dom does not encompass knowledge of bytes v. bits v. pits.

Would we be given this information on a test?

I feel like this is something we should definitely memorize by the end of a class. It could be a very useful statistic for those outside course 6 as well.

Also never mind my first comment. I just thought it looked hard because you said Electrical Engineering problem.

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

▶ *How far apart are the pits on a compact disc (CD) or CDROM?*

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

Its explained on the next page, but I initially was confused where the 10ˆ10 came from (rounding up 5*10ˆ9).

The factor of square root of two trick at the end is a pretty cool/easy way to make the answer so accurate... but not something I would have though of right away - it took a double take for me to see what had happened and why it worked.

I don't really understand this lattice idea, why does it need to spin then? does the spinning actually do anything for reading the disk?

I've never really heard of pits before. After reading this a couple times I think I understand how the distance was calculated, but it would help me to go over in class.

It was unclear to me that you were rounding 5x10ˆ9 to 10ˆ10.

It would be helpful to show how the math is done here too like right above just because this is the first example.

How would you know how big each pit is? Wouldn't that affect the spacing between?

I'm not certain, but I think that the idea that "each bit is stored in one pit" implies that their space size can be determined by the number of bits, which is in this case 5x10ˆ9. There is probably no extra-space (lattice), so the spacing may be the max space for the pit. I'm not sure that is says anything about the size of the actual pit though.

I think the idea is to get the distance from the "center" of one pit to the "center" of another pit, regardless of how big the pits are. There is a set # of pits in a finite space, so all we have to do is figure out how they are arranged relative to each other. The most useful way of describing their orientation relative to each other is by using their center position to the center of another pit.

The switch back to linear measure from the area is confusing. A simple illustration of the lattice and an enlargement of the pits might help.

I do not understand how the one pit per bit idea comes up. Is that an assumption? Or am I just missing a piece of fairly common information.

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

How far apart are the pits on a compact disc (CD) or CDROM?

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

In reality, are the pits arranged in a spiral track? This "lattice" approximation would be more like a square-like grid.

I had the same thought. On an actual CD the pits are arranged in a spiral so I originally thought that we would be looking at rings as opposed to boxes. The lattice approach is obviously more simple but is it okay for us to discount the way the pits are realistically arranged?

I'm curious how this is considered divide and conquer; it seems to me to be more of an estimation by rounding.

Is there any indication that the lattice should be square? Are the pits square and touching each other? Or is the radial spacing the same as the circumferential spacing? My first guess would be that the circumferential spacing would be small for faster reading, but the radial spacing larger so that you don't jump into the wrong row.

wouldn't this be the size of a pit plus the spacing between it and the next one?

I have the same question–maybe just for approximation purposes treat it as if there is no space between pits?

To me, at least in the context of our approximation, it seems that this quantity would be the size of the pit's "box" (if we envision our square CD as really small graph paper)

The comparison of the CD to a piece of graph paper makes sense to me, is it okay to envision it as so?

Good point. I should rephrase the original problem to be more clear. What I really meant, I now realize from reading your question and a couple others along the same lines, is the spacing between the *centers* of the pits – which is also what you described (size of pit plus spacing between the neighboring edges).

Is it generally okay to simplify circular things to rectangular?

I did the estimate using circular geometry and got .3 microns. I could be off, but I think the idea is that you can get within an order of magnitude fast. I'm curious as well

I think the estimate also becomes more valid given that the actual CD is larger 10 cm, but the middle portion doesn't contain pits

divide-and-conquer reasoning dissolves difficult problems into manageable pieces. It is a universal solvent for problems social, mathematical, engineering, and scientific.

To master any tool, try it out: See what it can do and how it works, and study the principles underlying its design. Here, the tool of divide and conquer is introduced using a mix of examples and theory. The three examples are CDROM design, oil imports, and the UNIX operating system; the two theoretical discussions explain how to make reliable estimates and how to represent divide-and-conquer reasoning graphically.

1.1 Example 1: CDROM design

The first example is from electrical engineering and information theory.

▶ *How far apart are the pits on a compact disc (CD) or CDROM?*

Divide finding the spacing into two subproblems: (1) estimating the CD's area and (2) estimating its data capacity. The area is roughly $(10\,\text{cm})^2$ because each side is roughly 10 cm long. The actual length, according to a nearby ruler, is 12 cm; so 10 cm is an underestimate. However, (1) the hole in the center reduces the disc's effective area; and (2) the disc is circular rather than square. So $(10\,\text{cm})^2$ is a reasonable and simple estimate of the disc's pitted area.

The data capacity, according to a nearby box of CDROM's, is 700 megabytes (MB). Each byte is 8 bits, so here is the capacity in bits:

$$700 \cdot 10^6\,\text{bytes} \times \frac{8\,\text{bits}}{1\,\text{byte}} \sim 5 \cdot 10^9\,\text{bits}.$$

Each bit is stored in one pit, so their spacing is a result of arranging them into a lattice that covers the $(10\,\text{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

$$d \sim \frac{10\,\text{cm}}{10^5} \sim 1\,\mu\text{m}.$$

It took me a few mins to figure out where these numbers came from ... it would be more useful to say: spacing CD length/# columns 10cm/10ˆ5 1um

Did Young's "double slit" experiment on this for a lab once. I think the answer came out to be something similar, so not too shabby on the estimate.

Are we going to get a units key on the test? Like how many micrometers are in a meter?

I think that as long as we stick with the metric system, it should be pretty straightforward. In any case, micro = 10ˆ-6, nano = 10ˆ-9, and kilo = 10ˆ3

I think that as long as we stick with the metric system, it should be pretty straightforward. In any case, micro = 10ˆ-6, nano = 10ˆ-9, and kilo = 10ˆ3

I'm having a hard time picturing this. Is this horizontal and vertical spacing or radial?

A diagram would be very helpful here!

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 μm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 μm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.
2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.
3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

I like the bit about adjusting the estimate by a factor of root 2. I do calculations like this often but don't adjust the results by the factors that I rounded with. This is a useful point and it could be worth emphasizing.

Doesn't the CD's shape come into play? It seems to be difficult to space them all the same considering the disc is circular. Or maybe not? I'm not sure. The estimate is still valid under the approximation that all pits are about the same distance away, I am just curious.

Also, the CD spins. Wouldn't that mean that if the pits were spaced the same, the ones farther out would have to be processed much faster than the ones near the center.

Is this why vinyl records have a higher sampling rate? It is an analog medium as opposed to a digital one, but if I recall correctly, my records are ripped via turntable to 24-bit tracks. Is this due to simply having more pits?

I am curious how the factor of 2 underestimates by sqrt(2)? Is it due to approximating the shape as a square?

> We are assuming the pits are arranged on the CD in a grid. To make the calculations easier, we assumed there were 10^10 total bits. However, this is off by about a factor of 2. Carrying this factor of 2 through the calculations, we get that there would actually be sqrt(10^10/2)=10^5/sqrt(2) rows, so d 10 cm/(10^5/sqrt(2)) = 10 cm / 10^5 *sqrt(2) 1 um * sqrt(2) = sqrt(2) um. So because we overestimated the number of bits by a factor of 2, we overestimated the number of rows by a factor of sqrt(2) (and also the # of col by sqrt(2)) and thus underestimated the distance by a factor of sqrt(2).
>
> Yes, I think you're on the right track. Evidence for approximating the shape as a square is at the bottom of page 4, where the text mentions that 10^10 pits would require 10^5 rows and 10^5 columns.
>
> > You are right, and it shows me that I should explicitly describe (and diagram) my approximate mental picture of a square lattice of pits.

I don't understand why we would do this simplification which only makes us do more work later. Why not plug in 5*10^9 bits in the first place?

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 μm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 μm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

Estimate the playing time, sampling rate, and sample size.

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.
2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.
3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

Is the factor of 2 important in doing an approximation? The difference in the final result is .4 micrometers, and even then that's also just an approximation.

I agree; how do we know what are appropriate approximations to make, and how do we know when the extra factors count?

This intuitively sounds right... It's about twice the wavelength of red light, and I'm pretty sure I've seen red lasers on various disk-reader modifications.

This all makes sense.

I feel like this is common knowledge now. I've always thought CDs were 700 MB, DVDs 4 GB, etc.

how are there cd's with different amounts of capacity? does that mean that the lattice is smaller?

what are the extremes of the size of storage and how is the price change relative. how do you make a cheap one vs a expensive one

this answers the previous question.

This sounds like a pretty specific way to estimate the CD's capacity. This is probably just a new way of thinking that I'm not used to, but I would consider this a calculation more than an estimation.

Some things are easily estimated, like the area of the CD, but I do agree that things like the sample rate aren't easily estimated, especially if you don't have too much knowledge on what the sampling rate is

I see it as several estimations used to calculate an estimate value...if each value that you are using in your formula is estimated, you're not really calculating the actual value

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 µm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 µm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.
2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.
3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

That's if you burn a CD, right? Does anyone know why many more (legally downloaded) mp3's can fit on a CD? Is that just data compression?

yes, mp3 files are compressed audio files. Traditional CD formats are very similar to the way that vinyl records worked, it's a 'note-for-note' style [vast over-simplification there]. where as mp3 files are smaller to store, but require more computing power to read.

Most industry produced CD's - like 2 disks "1 hits" - sit around this number per CD. Data compression shouldn't change it much.

I was also curious about this and went to wikipedia to read a little about CDs. In it it's stated that "Standard CDs have a diameter of 120 mm and can hold up to 80 minutes of uncompressed audio." So my assumption is that there is no data compression for CDs?

I thought CDs hold more like 80 minutes? 60 seems like a steep under-estimate.

I know right! I've been burning music to CD's for my car for years now. It's definitely 80 minutes (120 MB), but this is an estimation I guess. And 1 hour is a pretty number.

Not a big deal, but I think most CD's are 80 minutes?

That's a cool fact that I didn't know.

Interesting piece of info. Has the capacity remained constant?

Based on whether or not the capacity is 700mb vs. 650mb... the time varies between 74 and 80 minutes.

Are there actually several different accepted tempos to play a Beethoven symphony at?

Here is a challenge: We are assuming that we have never seen the box of a CD so we have never seen how many Megs are in a CD. Now, let's assume we have never listened to an Audio CD. Nowadays, we generally load CD's with Mp3s, which have compression, variable bit rate, and configurable frequency ranges... For what type of reader is this book intended? The derivation of unit conversions make it seem that this book is meant for high schoolers, but the assumption that the reader is familiar with audio formats makes it seem that the book is meant for readers with much more experience than average college students. The comments by other students in the class show that even MIT students can be unaware of digital audio formats...

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 µm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 µm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.

2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.

3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

I understand that with approximation, you use what you have to get to a reasonable estimate efficiently. Some numbers you don't have a feel for. Is it really possible to estimate anything?

I'm glad this is anonymous. Can someone explain what sampling rate means?

Sampling rate is the rate at which you sample the sound. To be more specific, a cd-rom only has a limited, finite amount of space for the pits. Therefore, we don't have the luxury of capturing a sound wave purely as the sinusoids that you see in calculus. Instead, we have to resort to picking points along the sinusoidal sound wave at regular intervals until we can make out the shape of the function (what its amplitude is, what its frequency is, etc.). Obviously, the more points you record, the more it will look like the original sound wave. The tradeoff occurs because too high a sampling rate will record an uncessary amount of points along the sound wave and waste precious memory space on the cd-rom

This really helped. Thanks

The explanation from a fellow student should help. And thanks for the comment. I should indeed explain, at least in passing, what sample rate means when I first use the term.

how am i supposed to know this?

Where do we get these numbers (20 kHz for hearing and the Nyquist Shannon theorem)? How do we proceed with the approximation when we don't know numbers like these.

My sense was that the first approximation, using just a ruler and what is on the CD box, is the one to use if you don't have any expertise. But it might be tempting, if you do have it (you know some information theory), not to try to approximate at all but to instead start trying to find the exact answer. Part of what I took away from the second approximation was: even if you have the expertise, you can use it to approximate, rather than just diving in trying to find the exact answer.

I also found the first approximation to be more intuitive. I had no idea what a sampling rate or sample size was so the second approximation was harder for me to comprehend. Is the second one supposed to be ideal for approximating or are we able to use whatever tactic works best for us personally?

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 µm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 µm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.
2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.
3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

Does the ability to estimate depend on this type of knowledge?

I would not expect most readers (not even engineers) to be very familiar with wave file formats. Are you using these figures only to illustrate unit conversions or are you telling the reader that he/she should know all of this common sense? Many readers might feel intimidated by this section, if it isn't stated that the section is just meant to illustrate unit conversions.

I would have been unable to come up with any of these estimates myself due to lack of knowledge on this field. Is there any other way I could have come up with these estimates?

I agree, this is a little beyond common knowledge.

I don't quite understand how sampling rate directly can be correlated to the frequency at which we can hear sound. I am assuming that the disc reader gathers some information from the CD stores it and then plays it. I don't think it plays directly as fast as it reads. Thus the acoustics analogy seems to not quite work. Pits can only relay a yes or no respone (1 or 0). How can it dictate volume, and key which cover a large range? Maybe I don't know enough about acoustics and how sound and cd players work.

what does a telephone do to the frequency of sound? and why does it change it? is there a way to get a telephone that does not skip at high frequencies

What does the 44kHz sampling rate refer to?

I think it refers to how often the disc reader scans for bits or no bits. 44khz means it scans 44,000 times per second. I think.....

the best way to explain it is with pictures: http://en.wikipedia.org/wiki/Sampling_rate

Where can we learn more about this theorem?

Explain theory?

I think when sampling a signal you need to sample it at least twice every period, ideally once at the top of the wave and once at the bottom. Otherwise it can look like a way slower frequency because you are losing so much data by sampling infrequently.

Otherwise, you get aliasing and data loss

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 μm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 μm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.
2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.
3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

What is sampling?

Its a way of turning a continuous wave (like a sine wave, and the sound that would result) into discrete parts. You want a high enough sampling frequency to pick up on each wave at least twice. Basically, its chopping a wave into tiny little pieces that to us would not sound any different that the entire wave, only now it can be read out in bits!

How do high rates simplify the antialias filter?

And what is an antialias filter? What does it do?

> An antialias filter is a device (or, often, a program) that prevents aliasing. But this is not a helpful answer on its own; perhaps the background question is, "What is aliasing?" I'll explain that in lecture on Friday, with a demonstration if I can figure out a good one.

For a decent review, check out pages 6-17 of Lecture 19: https://stellar.mit.edu/S/course/20/

(courtesy of 20.309 last semester)

So these 'standards' for rates are designed so that in total the songs, when translated into digital data, occupy the full (or approximately) 700MB using up the pits on the disc surface? I know that lossless files, being exact rips off the CD, can vary greatly in actual bitrate from between 700 to even upwards of 1200 kbps.

How does one obtain these seemingly random yet useful facts? And then how do you know to apply it

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 µm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 µm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.

2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.

3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

this doesn't really help me to understand this concept.

ditto. i don't have any background in signals, etc., and found this section very difficult to follow.

my guess: you want the highest rate possible, given that it must be above 40kHz to make sure you actually get everything (no aliasing). The higher, the better. So they wanted to try 80kHz, but either that cost too much to make or computers could not write to the CD effectively enough with that many pits. So, they went with 40kHz, and assumed that there would be an error margin of around 4kHz. Thus, the set sampling rate became 44kHz, to allow room for that 4kHz margin

This was difficult for me as well and felt a little out of place given the straightforwardness of other examples.

I am also unclear about this.

I also found this section very difficult to understand - given how little specialized knowledge the other examples required, it also felt out of place.

Me too..I don't understand the sampling rate part.

What does the margin mean?

If we didn't know these numbers and the Nyquist-Shannon theorem, is there any other way to proceed?

How are sampling margin and sampling rates related?

It still isn't clear to me the difference between sampling rate margin and sampling rate. How can sampling rate margin be 4kHz and the sampling rate be 44kHz, when reconstructing a 20kHz signal only requires sampling at 40kHz?

I think the sampling rate margin is how much faster he's sampling than what's required. 40kHz is required but he's sampling at 44kHz giving a difference, or margin, of 4kHz.

I'm a little confused on your use of the word utopia.

What exactly is a sample defined as? A sub-set of bits/pits?

That calculation was simplified by rounding up the number of bits from $5 \cdot 10^9$ to 10^{10}. The factor of 2 increase means that 1 μm underestimates the spacing by a factor of $\sqrt{2}$, which is roughly 1.4: The estimated spacing is 1.4 μm.

Finding the capacity on a box of CDROM's was a stroke of luck. But fortune favors the prepared mind. To prepare the mind, here is a divide-and-conquer estimate for the capacity of a CDROM – or of an audio CD, because data and audio discs differ only in how we interpret the information. An audio CD's capacity can be estimated from three quantities: the playing time, the sampling rate, and the sample size (number of bits per sample).

▶ *Estimate the playing time, sampling rate, and sample size.*

Here are estimates for the three quantities:

1. *Playing time.* A typical CD holds about 20 popular-music songs each lasting 3 minutes, so it plays for about 1 hour. Confirming this estimate is the following piece of history. Legend, or urban legend, says that the designers of the CD ensured that it could record Beethoven's Ninth Symphony. At most tempos, the symphony lasts 70 minutes.
2. *Sampling rate.* I remember the rate: 44 kHz. This number can be made plausible using information theory and acoustics.

   First, acoustics. Our ears can hear frequencies up to 20 kHz (slightly higher in youth, slightly lower in old age). To reproduce audible sounds with high fidelity, the audio CD is designed to store frequencies up to 20 kHz: Why ensure that Beethoven's Ninth Symphony can be recorded if, by skimping on the high frequencies, it sounds like was played through a telephone line?

   Second, information theory. Its fundamental theorem, the Nyquist–Shannon sampling theorem, says that reconstructing a 20 kHz signal requires sampling at 40 kHz – or higher. High rates simplify the antialias filter, an essential part of the CD recording system. However, even an 80 kHz sampling rate exceeded the speed of inexpensive electronics when the CD was designed. As a compromise, the sampling-rate margin was set at 4 kHz, giving a sampling rate of 44 kHz.
3. *Sample size.* Each sample requires 32 bits: two channels (stereo) each needing 16 bits per sample. Sixteen bits per sample is a compromise between the utopia of exact volume encoding (infinity bits per sample

How many people know this? I don't feel like this is necessarily common knowledge. But perhaps it is just me.

I agree, I'm not familiar with this information. I feel like this section lost its meaning for me because I didn't know any of this information.

Agreed. Most of the other numbers have been familiar. This one (16 bits per sample) I have no idea where it came from.

Same here. What is sample size (as someone else asked) and how the bits related seemed unclear to me.

This whole section makes a lot more sense if you've taken 6.003 or its equivalent. Perhaps it would be better to introduce the section with a short description of the physical structure of a CD and how sound is taken from it.

per channel) and the utopia of minimal storage (1 bit per sample per channel). Why compromise at 16 bits rather than, say, 50 bits? Because those bits would be wasted unless the analog components were accurate to 1 part in 2^{50}. Whereas using 16 bits requires an accuracy of only 1 part in 2^{16} (roughly 10^5) – attainable with reasonably priced electronics.

The preceding three estimates – for playing time, sampling rate, and sample size – combine to give the following estimate:

$$\text{capacity} \sim 1\,\text{hr} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4\times 10^4\,\text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\,\text{sample}}.$$

This calculation is an example of a conversion. The starting point is the 1 hr playing time. It is converted into the number of bits stepwise. Each step is a multiplication by unity – in a convenient form. For example, the first form of unity is 3600 s/1 hr; in other words, 3600 s = 1 hr. This equivalence is a truth generally acknowledged. Whereas a particular truth is the second factor of unity, $4.4\cdot 10^4$ samples/1 s, because the equivalence between 1 s and $4.4\cdot 10^4$ samples is particular to this example.

Problem 1.1 General or particular?
In the conversion from playing time to bits, is the third factor a general or particular form of unity?

Problem 1.2 US energy usage
In 2005, the US economy used 100 quads. One quad is one quadrillion (10^{15}) British thermal units (BTU's); one BTU is the amount of energy required to raise the temperature of one pound of liquid water by one degree Fahrenheit. Using that information, convert the US energy usage stepwise into familiar units such as kilowatt–hours.

What is the corresponding power consumption (in Watts)?

To evaluate the capacity product in your head, divide it into two subproblems – the power of ten and everything else:

1. *Powers of ten.* They are, in most estimates, the big contributor; so, I always handle powers of ten first. There are eight of them: The factor of 3600 contributes three powers of ten; the 4.4×10^4 contributes four; and the 2×16 contributes one.

Comments on page 4

I was only able to understand this after reading it 2 or 3 times.

What exactly is meant by "accurate to 1 part in 2^X"?

I think it means that 50 bits per sample is too accurate–like it doesnt make much of a difference (or the analog components aren't that accurate anyway), so a smaller number like 16 bits per sample is more commonly used

I agree - without exceedingly accurate components that could reproduce the waveform exactly as instructed by the media (i.e., cd), the higher sample rate would just be down-sampled anyway. Perhaps this accounts for the perception that vinyl sounds "better" - it isn't limited in sample rate?

As an example of the same language, a resistor whose resistance is accurate to 1 part in 20 means that the resistance may be incorrect by 1/20 or 5%.

So "accurate to 1 part in 2^50" means accurate to 1/2^50, which is 10^-15 (remind me in lecture to show everyone how to do that calculation mentally). In other words, the analog hardware would have to be precise in its specifications to 15 decimal places. That is difficult to achieve, to understate the problem; it is analogous to measuring the distance to the sun to within 0.1 mm.

at this point, the section stops being about estimating playing time, sampling rate, & sampling size...which made it a little hard to follow, since it's still under that heading.

This whole section is really about estimating the capacity by breaking it down...it would make more sense to have the title of the section be 'Estimating the capacity from playing time, sampling rate, and sample size.' and maybe put the heading above the "finding the capacity on a box..." paragraph

...it might make it flow better for those of us (maybe it was just me) who don't have time to read it all at once and pause at the mini-headings.

This makes complete sense, but maybe you could put this explanation earlier when you use conversion for the first time.

per channel) and the utopia of minimal storage (1 bit per sample per channel). Why compromise at 16 bits rather than, say, 50 bits? Because those bits would be wasted unless the analog components were accurate to 1 part in 2^{50}. Whereas using 16 bits requires an accuracy of only 1 part in 2^{16} (roughly 10^5) – attainable with reasonably priced electronics.

The preceding three estimates – for playing time, sampling rate, and sample size – combine to give the following estimate:

$$\text{capacity} \sim 1\,\text{hr} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4 \times 10^4\,\text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\,\text{sample}}.$$

This calculation is an example of a conversion. The starting point is the 1 hr playing time. It is converted into the number of bits stepwise. Each step is a multiplication by unity – in a convenient form. For example, the first form of unity is 3600 s/1 hr; in other words, 3600 s = 1 hr. This equivalence is a truth generally acknowledged. Whereas a particular truth is the second factor of unity, $4.4 \cdot 10^4$ samples/1 s, because the equivalence between 1 s and $4.4 \cdot 10^4$ samples is particular to this example.

Problem 1.1 General or particular?

In the conversion from playing time to bits, is the third factor a general or particular form of unity?

Problem 1.2 US energy usage

In 2005, the US economy used 100 quads. One quad is one quadrillion (10^{15}) British thermal units (BTU's); one BTU is the amount of energy required to raise the temperature of one pound of liquid water by one degree Fahrenheit. Using that information, convert the US energy usage stepwise into familiar units such as kilowatt–hours.

What is the corresponding power consumption (in Watts)?

To evaluate the capacity product in your head, divide it into two subproblems – the power of ten and everything else:

1. *Powers of ten.* They are, in most estimates, the big contributor; so, I always handle powers of ten first. There are eight of them: The factor of 3600 contributes three powers of ten; the 4.4×10^4 contributes four; and the 2×16 contributes one.

Great method of converting units. I learned it in 7th grade, and still find it immensely useful. It's nice that you included the equation rather than glossing over it as some textbooks do.

Often times the written descriptions are hard to follow and being a very visual learner, I find these dimensional conversions extremely helpful for comprehension and memorization.

This also helps make sure the units come out the way they should.

dimensional analysis

I don't think they're the same thing, though I can't articulate it now.

I think dimensional analysis is a tool to determine relationships given different important variables and their dimensions in terms of unitless length, time, weight, etc.. This is a conversion because it starts with capacity = 1hr, and then converts 1 hr into bits.

I agree with that - I think that dimensional analysis is being used, but only to inform the process of converting one hour into bits.

[from the same poster as at 5:41] Err, to be clear, I don't think any dimensional analysis is being used at all in this example.

see: http://en.wikipedia.org/wiki/Dimensional_analysis#Examples

I don't actually know what multiplication by unity means. All the factors in this example are a ratio of something to one. Is that what makes it "by unity"?

Don't be misled by the "1" in "1 hr", in "1 s", or in "1 sample". What I meant (and will try to clarify when i revise this spot) is that 1 hr = 3600 s, so (3600 s)/(1 hr) is, by definition, equal to 1: You are dividing two identical quantities.

And multiplying by "1" (or "unity") can never hurt you. If you use a suitable ratio, then you can convert from one unit system to another (here, from hours to bits, in three steps of multiplication by unity).

So...dimensional analysis, right?

This is a very basic, necessary step and could probably be explained in simpler terms.

per channel) and the utopia of minimal storage (1 bit per sample per channel). Why compromise at 16 bits rather than, say, 50 bits? Because those bits would be wasted unless the analog components were accurate to 1 part in 2^{50}. Whereas using 16 bits requires an accuracy of only 1 part in 2^{16} (roughly 10^5) – attainable with reasonably priced electronics.

The preceding three estimates – for playing time, sampling rate, and sample size – combine to give the following estimate:

$$\text{capacity} \sim 1\,\text{hr} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4 \times 10^4\ \text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\,\text{sample}}.$$

This calculation is an example of a conversion. The starting point is the 1 hr playing time. It is converted into the number of bits stepwise. Each step is a multiplication by unity – in a convenient form. For example, the first form of unity is 3600 s/1 hr; in other words, 3600 s = 1 hr. This equivalence is a truth generally acknowledged. Whereas a particular truth is the second factor of unity, $4.4 \cdot 10^4$ samples/1 s, because the equivalence between 1 s and $4.4 \cdot 10^4$ samples is particular to this example.

Problem 1.1 General or particular?

In the conversion from playing time to bits, is the third factor a general or particular form of unity?

Problem 1.2 US energy usage

In 2005, the US economy used 100 quads. One quad is one quadrillion (10^{15}) British thermal units (BTU's); one BTU is the amount of energy required to raise the temperature of one pound of liquid water by one degree Fahrenheit. Using that information, convert the US energy usage stepwise into familiar units such as kilowatt–hours.

What is the corresponding power consumption (in Watts)?

To evaluate the capacity product in your head, divide it into two subproblems – the power of ten and everything else:

1. *Powers of ten.* They are, in most estimates, the big contributor; so, I always handle powers of ten first. There are eight of them: The factor of 3600 contributes three powers of ten; the 4.4×10^4 contributes four; and the 2×16 contributes one.

This language and syntax gets in the way of meaning in explaining conversion, which is a really important and misunderstood skill.

I understand there are 3600 sec in an hour, but where do 4.4*10^4 come from? Why is the sampling rate used as a conversion unit?

I think 4.4*10^4 samples comes from the 44kHz from the previous page.

I both agree and disagree. I think there are some extraneous words and this makes these couple of sentences more difficult to understand than necessary. In general though I like the explanation of "forms of unity" in changing one measured unit into another equivalent unit.

On the language and syntax (postposition adjective in "truth generally acknowledged"): It was an allusion to the opening line of a famous novel: "It is a truth universally acknowledged, that a single man in possession of a good fortune, must be in want of a wife." But I misremembered it and wrote "generally acknowledged" instead of "universally acknowledged". I'll try to clarify here.

I have no idea what this problem is asking. What does it mean when it asks if something is "general" or "particular"?

I don't understand this question.

yeah I agree, this sentence is confusing...I'm not sure what the last part of the question means.

1.1) The factor is particular, due to the definition of a sample in this example.

I'm not sure what Problem 1.1 is asking?

per channel) and the utopia of minimal storage (1 bit per sample per channel). Why compromise at 16 bits rather than, say, 50 bits? Because those bits would be wasted unless the analog components were accurate to 1 part in 2^{50}. Whereas using 16 bits requires an accuracy of only 1 part in 2^{16} (roughly 10^5) – attainable with reasonably priced electronics.

The preceding three estimates – for playing time, sampling rate, and sample size – combine to give the following estimate:

$$\text{capacity} \sim 1\,\text{hr} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4 \times 10^4\,\text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\,\text{sample}}.$$

This calculation is an example of a conversion. The starting point is the 1 hr playing time. It is converted into the number of bits stepwise. Each step is a multiplication by unity – in a convenient form. For example, the first form of unity is 3600 s/1 hr; in other words, 3600 s = 1 hr. This equivalence is a truth generally acknowledged. Whereas a particular truth is the second factor of unity, $4.4 \cdot 10^4$ samples/1 s, because the equivalence between 1 s and $4.4 \cdot 10^4$ samples is particular to this example.

Problem 1.1 General or particular?

In the conversion from playing time to bits, is the third factor a general or particular form of unity?

Problem 1.2 US energy usage

In 2005, the US economy used 100 quads. One quad is one quadrillion (10^{15}) British thermal units (BTU's); one BTU is the amount of energy required to raise the temperature of one pound of liquid water by one degree Fahrenheit. Using that information, convert the US energy usage stepwise into familiar units such as kilowatt–hours.

What is the corresponding power consumption (in Watts)?

To evaluate the capacity product in your head, divide it into two subproblems – the power of ten and everything else:

1. *Powers of ten.* They are, in most estimates, the big contributor; so, I always handle powers of ten first. There are eight of them: The factor of 3600 contributes three powers of ten; the 4.4×10^4 contributes four; and the 2×16 contributes one.

The third is a particular truth. But i feel like here is where a lot could go wrong. pretty much everyone knows that a CD is about an hour long, slightly longer. But for me at least, it doesn't seem like common knowledge that our ears hear frequencies up to 20kHz so the sampling rate should be 44kHz; and that there are 16 bits for each of two channels. slight mistakes here could cause our estimation to easily be 2 or 4 times the real value

I have a similar concern - it seems like much of the estimation here relies on preexisting knowledge. How would we make such an estimate without this knowledge? Do we always just start with what we know and try to apply it to the problem at hand?

I agree it's frustrating when you don't have any idea how large or small something is. What do you do then?

I agree. It's definitely a particular truth.

Approximation problems like this are always rely on knowing a few specific numbers, which is why it's usually helpful to work in groups. While you may not know a certain number, someone else in your group can. Through my experience, it seems like a group of 4 or 5 people can come pretty close to a correct approximation.

I think it's particular because its based on the sampling frequency we estimated earlier.

I don't know what "general" and "particular" forms of unity are. I have never heard of them.

I think a "general" form of unity is an equivalence that is generally accepted to be true like 12 inches in a foot or 10 centimeters in a meter. A "particular" form of unity is something you yourself has set to be true for the unique problem your working on. Or its a pretty good assumption that you can make to solve a bigger problem.

particular. right?

this interrupts the thought process started above. it should be moved.

Problem 1.1 is a little random and doesn't really have any continuity with what the passage is trying to explain.

An interesting side note to the 44Khz number (our bank of numbers) it can be factored as 2^2 * 3^2 * 5^2 * 7^ 2....which happens to be a product of the squares of the first 4 prime numbers...coincidence?

per channel) and the utopia of minimal storage (1 bit per sample per channel). Why compromise at 16 bits rather than, say, 50 bits? Because those bits would be wasted unless the analog components were accurate to 1 part in 2^{50}. Whereas using 16 bits requires an accuracy of only 1 part in 2^{16} (roughly 10^5) – attainable with reasonably priced electronics.

The preceding three estimates – for playing time, sampling rate, and sample size – combine to give the following estimate:

$$\text{capacity} \sim 1\,\text{hr} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4\times10^4\ \text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\ \text{sample}}.$$

This calculation is an example of a conversion. The starting point is the 1 hr playing time. It is converted into the number of bits stepwise. Each step is a multiplication by unity – in a convenient form. For example, the first form of unity is 3600 s/1 hr; in other words, 3600 s = 1 hr. This equivalence is a truth generally acknowledged. Whereas a particular truth is the second factor of unity, $4.4\cdot10^4$ samples/1 s, because the equivalence between 1 s and $4.4\cdot10^4$ samples is particular to this example.

Problem 1.1 General or particular?
In the conversion from playing time to bits, is the third factor a general or particular form of unity?

Problem 1.2 US energy usage
In 2005, the US economy used 100 quads. One quad is one quadrillion (10^{15}) British thermal units (BTU's); one BTU is the amount of energy required to raise the temperature of one pound of liquid water by one degree Fahrenheit. Using that information, convert the US energy usage stepwise into familiar units such as kilowatt–hours.

What is the corresponding power consumption (in Watts)?

To evaluate the capacity product in your head, divide it into two subproblems – the power of ten and everything else:

1. *Powers of ten.* They are, in most estimates, the big contributor; so, I always handle powers of ten first. There are eight of them: The factor of 3600 contributes three powers of ten; the 4.4×10^4 contributes four; and the 2×16 contributes one.

It's very hard to distinguish comments on a particular section when everybody highlights only the margin of the text. It forces people to slowly make their comments wider and wider.

What is this doing here?

10^15 BTUs * (0.454 kg/ 1 lb) * (5 deg C/ 9 deg F) * (4186 J/kg/deg C) = Joules

15 - 1 - 1 + 3 = 16 zeros 4.54*5.55*4.186 = few * 10 * few = 100

1*10^18 Joules

It's a little hard to tell because of the formatting issues, but I think this conversion is of 1 quad to Joules, instead of 100 quads. I think the result should be about 10^20 Joules (or about 3*10^13 kWh)

1.2) (1000 quad/yr) * (10^15 BTU/quad) * (1 W*hr/3 BTU) * (1 yr/9000 hr) -> 3*10^13 W Power consumption over the year.

This corresponds to 3*10^14 kW-hrs of energy.

The 1st approximation is a slight underestimate of BTU to W*hr, but the 2nd approximation is a slight overestimate of hr to yr. So, things should about work out.

specific heat of water = 4.18 J/gC

100 quads * (10^15 BTUs / 1 quad) * (1 kWh / 3400 BTUs) 1/34*10^15 1/3*10^14 3*10^13

10^18 Joules / 1 year * (1 year/ 365 days) * (1 day / 24 hours) * (1 hr / 60 mins) * (1 min / 60 sec)

18 - 3 - 2 - 2 - 2 = 9 zeros 2.73*4.16*1.66*1.66 = few*10*1*1 = 30

3*10^10 Watts (Joules/sec)

Looks good to me

could/should this be another heading...estimating complicated math?

Do you usually do these problems in your head? or on paper?

Interesting idea, I haven't broken up a math problem like that before.

per channel) and the utopia of minimal storage (1 bit per sample per channel). Why compromise at 16 bits rather than, say, 50 bits? Because those bits would be wasted unless the analog components were accurate to 1 part in 2^{50}. Whereas using 16 bits requires an accuracy of only 1 part in 2^{16} (roughly 10^5) – attainable with reasonably priced electronics.

The preceding three estimates – for playing time, sampling rate, and sample size – combine to give the following estimate:

$$\text{capacity} \sim 1\,\text{hr} \times \frac{3600\,\text{s}}{1\,\text{hr}} \times \frac{4.4 \times 10^4\,\text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\,\text{sample}}.$$

This calculation is an example of a conversion. The starting point is the 1 hr playing time. It is converted into the number of bits stepwise. Each step is a multiplication by unity – in a convenient form. For example, the first form of unity is 3600 s/1 hr; in other words, 3600 s = 1 hr. This equivalence is a truth generally acknowledged. Whereas a particular truth is the second factor of unity, $4.4 \cdot 10^4$ samples/1 s, because the equivalence between 1 s and $4.4 \cdot 10^4$ samples is particular to this example.

Problem 1.1 General or particular?
In the conversion from playing time to bits, is the third factor a general or particular form of unity?

Problem 1.2 US energy usage
In 2005, the US economy used 100 quads. One quad is one quadrillion (10^{15}) British thermal units (BTU's); one BTU is the amount of energy required to raise the temperature of one pound of liquid water by one degree Fahrenheit. Using that information, convert the US energy usage stepwise into familiar units such as kilowatt–hours.

What is the corresponding power consumption (in Watts)?

To evaluate the capacity product in your head, divide it into two subproblems – the power of ten and everything else:

1. *Powers of ten.* They are, in most estimates, the big contributor; so, I always handle powers of ten first. There are eight of them: The factor of 3600 contributes three powers of ten; the 4.4×10^4 contributes four; and the 2×16 contributes one.

This is definitely a useful technique. I use this all the time for estimation and to make calculations faster.

Why not call it 32 as in the expression above?

Because we are focused on evaluating powers of 10. The 2 X 16 more clearly shows a power of ten in this number.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6 \times 4.4 \times 3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6 \times 4.4 \times 3.2$ is approximately $(\text{few})^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5 \cdot 10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

Comments on page 5

When there is a huge box of highlights, how do I read the smaller boxes within the huge box?

That's an awesome word.

I never knew that was a word for this, thats cool

Is this the actual term for any digits less than 10?

does this mean if we were doing 6ˆ3 we would be doing a fewˆ3 still? how does this work for other numbers? doing 3x4x3 was very convenient. Also, I can't see the other comments for this page because something is probably wrong with my computer so I apologize if this question is repeated.

I'm confused... halfway between 1 and 10 is 5, but 'few' is defined at 3? Why not just round all numbers greater than 5 to 10 and all less to 1?

There has to be a better way to explain why 3 is the critical number to look for. Perhaps instead of introducing it as "halfway" between 1 and 10, you say something like: "we're looking for factors that, when multiplied, increase the answer by an order of magnitude. Therefore we define "few" to be the number that when squared is 10 - that number is about 3.

Isn't halfway between 1 and 10 "5.5"? But then you say that few=3. Do you just mean between 1 and 10? not 'halfway'?

He describes what he means by "few" in paragraph 2: "'few' is halfway between 1 and 10; therefore...(few)ˆ2=10" so few 3, not five. I agree however that saying "'few' is halfway between 1 and 10" definitely implies five, so perhaps it should read "'few' is between 1 and 10 such that (few)ˆ2 = 10"?

I don't entirely understand how "few" is being defined here (and how it is being used)... if there were more values being multiplied, how would this work?

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6\times4.4\times3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6\times4.4\times3.2$, each factor is roughly a 'few', so $3.6\times4.4\times3.2$ is approximately $(\text{few})^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5\cdot10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

Just curious, but is this a well known definition? I've heard "few" used to mean a value >3 (and sometimes >5), but haven't seen it expressed like this? When explaining an approximation to someone outside this class would "few" be generally accepted to mean this?

I had certainly never heard it before this class, but it does make sense in the context of accuracy to within a power of 10.

That's good to know. Was this the original idea behind "few"? And how does one know how much to overestimate by? if the number were closer to 9, then it would automatically be few^2 already...

If "few" is halfway in between 1 and 10, isn't it 5? And doesn't 5^2=25?

I don't fully understand the concept of the 'few'–is it always 3, or is it only 3 in this example (because the numbers are all below 5)? Because it says 'few' is halfway between 1 and 10 which would be 5, but it also says (few)^2=10, so few 3

Ah, good point. "Few" is always 3, or more exactly sqrt(10), which is about 3.16. "Few" is indeed halfway between 1 and 10, but on a logarithmic scale: Going from 1 to few should be the same factor as going from few to 10. That's why few is sqrt(10) or, for simplicity, about 3.

Hi Sanjoy, I'm testing the direct reply feature !

This is super useful to know, and I wish I had been taught this trick earlier in school.

What do you do with numbers that aren't around 3?

It would be nice to have a box under the "everything else" section with a re-written version of the "capacity ..." line from the previous page. It could visually sum up the approximation computation you've just described. Eg it could read "capacity (10^3 x 10^4 x 10) x (few x few x few) = 10^9 x 50", or something like that.

Would it be inefficient to just round to the nearest whole number and multiply out that way? For example, this would turn into 4 x 4 x 3 = 48, which was close to the guess of 50 and and only a little further from the real answer of 50.688.

Yeah, I agree.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6 \times 4.4 \times 3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6 \times 4.4 \times 3.2$ is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5 \cdot 10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

I've definitely never heard of this logic with "few" before, and although I think I get it from this reading, I would prefer to talk more about it in class.

Does it matter whether we express the numbers only in terms of exponents of few or if we use exponents of 10 multiplied by few.

The idea of "few" is really cool, never heard of that before.

I would like to know a bit more about when this type of approximation is alright to use. Clearly, there are problems which require a more specific approximation of the measurements.

What is an acceptable amount of error in these cases? 50% 25%?

If you used 'few'=5 here, the number would be way off–almost two orders of magnitude off, just from picking a slightly different 'few'. how do you choose what your 'few' is?

Where is the thought process to add 20?

I understand why 30 is an underestimate, but what I don't understand is why 50 is the next number guessed, it's almost double the original. Is it because it's easier to work with?

How do you know to guess 50 from getting (few)^3=30?

How do we figure out what the correction factor should be?

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6\times4.4\times3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6\times4.4\times3.2$ is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5{\cdot}10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

To be honest, by the time I read through this section of the paper, I forgot its original purpose was to discuss divide and conquer. I just feel like there are much better examples out there for divide and conquer than what was discussed.

I agree. I tend to think of divide & conquer as breaking the problem into smaller sub-problems, but this felt more like dimensional analysis. Could you maybe clarify the exact definition of divide & conquer as it pertained to these examples?

I think that dimensional analysis is a good example of divide and conquer. You have to look at little pieces of things and put them together...like part 2, to find the capacity (ie bits) of a 'typical' cd you first need to think about how digital audio is encoded [in samples, you take something analog and give it discreet steps...like pixels in an image]. how many 'audio samples' make up a CD? well, how fast are you sampling it (rate)? how many songs at what length (time)? how big is a sample (size)?

yes, you are using simple dimensional analysis when you put it all together, but to get to that point you need to divide up the problem into parts that can be analyzed.

So in summary, the divide and conquer method is a procedure in which you can branch out the different parts of a problem to facilitate its solution?

I can see how Divide and Conquer is a powerful method, so I feel that it should only be used in complex problems. Otherwise you would probably be wasting time with this. And I feel like everyone uses this method unconsciously.

Most interesting part of it all.

I was confused in the explanation of this section.

Actually, after being really confused about the explanation, I reread the last paragraph on page 6 and the first one on page 7...now I understand the technique of dividing the calculation into powers of ten and everything else. I think what initially threw me off was the word mantissas and because I didn't realize that section was still talking about the calculation above. Maybe not having the grey box between those two sections would have helped.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product 3.6×4.4×3.2. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6 \times 4.4 \times 3.2$ is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5 \cdot 10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

If I did not know how to solve this problem myself, it might be hard to see how the whole section illustrated the divide-and-conquer approach. The section was more focused on approximation methods, unit conversions, and solving the problem than in showing the modularity of the divide and conquer strategy. Diagrams would be very helpful to show how each "module" relates to another and how they all work together to answer the question. I still do not clearly see the point of deriving numbers in so many different ways. Why assume that I do not know how many Mb's of data a CD has, but that I do know the bitrate of a CD? Perhaps the theme that approximations should agree with one another and with the real answer is not being explicitly stated here.

Once we receive an estimate using divide and conquer, what would be the best way to verify our result, if at all? Will using a second divide and conquer approach be sufficient?

I'm glad you asked! For this problem, the answer will be in the next reading (section 1.2), which will be posted on Friday (the memo will be due on Sunday at 10pm).

I feel like this is something I tend to do a lot of in day-to-day life...I'm looking forward to reading the next section...finesse is a good thing!

I feel a lot of times you can just compare the estimate to the actual answer. If it matches or almost matches, you did a good job! If you got something in the hundreds when you were suppose to get something in the millions, that tells you something is wrong.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6\times4.4\times3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6\times4.4\times3.2$ is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5{\cdot}10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

How does this relate to divide and conquer other than checking solutions we found that way?

I'm not sure either. I think the point of this article is to show that checking your answer with different methods is best though. Also, these examples were confusing to me, so I might be missing something.

It relates to divide and conquer because the redundant solutions use divide and conquer, but also to address an inherent flaw in the divide and conquer method. When we break down a problem in a particular way, we may inherently bias our answer (we pretended the CD was a square). Without checking using a different route, our confidence in our estimate is limited.

I am amazed by how you can come up with such methods to calculate the spacing? How do you know that you can utilize the diffraction principles etc, I would have never thought of that? Will I be able to relate problems as such to the different principles I have learned from different classes by the end of the semester?

I agree, I think the first and hardest step is actually realizing that one may actually have knowledge that is applicable to help solve an approximation problem. I would never have thought to think about the CD player to predict spacing.

Is there some sort of method that you can default to if you do not know an equation? Just multiply them maybe?

I don't think there will ever be a default to fall back on, but as others said I'm realizing that I do have knowledge to help me estimate. However I feel like to apply knowledge such as wave diffraction one would need to conduct a small test. Easy for a CD, hard for a 747.

This experiment was done in 8.02 and something similar was done in 5.112 – makes it easier for me to visualize.

Why does the laser wavelength need to be smaller than the pit spacing?

The wavelength is tuned to the distance of where the laser is emitted and the bottom of the pit on the CD so that when the light hits a pit, destructive interference occurs. If the wavelength was larger than the pit, it would be impossible for the reflected light to be knocked out of phase by half a wavelength.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6\times4.4\times3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6\times4.4\times3.2$ is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5{\cdot}10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

Does "divide and conquer" as used in this paragraph refer to an umbrella approach to answering any question? As in, divide the way one approaches a given problem and solve it in different ways in order to check your estimation? I agree with previous comments, this paragraph is a bit nebulous.

I found what might be a bug with the NB script. When I already have one of the files loaded, and I click the "Files" tab and select a document, the script becomes unresponsive and crashes Firefox. It works fine when I go to the NB in the beginning with no particular document loaded though. Also, NB doesn't work in Safari at all.

Another bug: When typing a comment, if you try to click to another part of the comment to edit something (say a spelling error) the drop down just reappears there and I cannot make the edit. This is only the case for notes on the right side; not globals.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6 \times 4.4 \times 3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6 \times 4.4 \times 3.2$ is approximately $(\text{few})^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5 \cdot 10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

I just learned a new word! Are there lots of terms used like this in the text?

I did the homework before reading this. Now your solutions make a lot of sense when you say few.

I think if you were to publish a new edition to this book, you should mention that since we are dealing with order of magnitude estimation, 10ˆ.5 is about 3 is a justified reason to claim few=3. When you explained this in class, it really helped me feel comfortable with applying that approximation.

I agree with this. Maybe mentioning that the exponents are related to keeping the estimate in a logarithmic scale would be helpful.

This is one instance where an example is amazing. I would prefer to have it written in equation form, outside the body of the paragraph, to show exactly how a calculation would resolve into fewˆsome power * 10ˆsome power. This particular example may not be best in this case.

Why do they say to calculate (few)ˆ3 = 30, if they change their answer to 50?

I think it's done just cuz 50 is an easier number to work with than 30.

He explains in the next line that because each number is slightly greater than 3, the few technique is a bit of an underestimate so he bumps it up to 50.

how do you know to guess 50, is this a gut feeling problem again? I know it would be greater than 30, but I don't really know whether to guess 40 or 50.

I think the general theme is that it doesn't really matter whether you guess 40 or 50 – they're pretty close to each other, and we're only aiming for an estimate anyway. One way to guess though, is that 3.2ˆ2 10, and 3.6 is roughly 10% greater than 3.2 (actually 12.5%). So 3.6*4.4*3.2 3.2ˆ2*1.1*4.4 3.2ˆ2*4.8 10 * 5.

Regarding to the notes visibility button: It would be nice to have an option to filter out notes by person or to just look at one's notes.

This comment is exactly what I was going to post. A little bit back I tried to look at my last post and was wondering how to do that. It seems there is no way... is there?

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6\times4.4\times3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6\times4.4\times3.2$ is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5\cdot10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

Find the examples of divide-and-conquer reasoning in this section.

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

I'm not sure how you could fix this, but when I started reading section 1.2, I assumed that we would be presented with other divide and conquer methods, like we saw in section 1.1. However, this section is about ways to ensure the accuracy of a divide and conquer approximation. This section just seemed out of place since only the one example of the CD has been given so far.

Read Section 1.2 and do the memo by Sunday at 10pm. Have fun!

If it can make the reading easier, check out this **untested pre-alpha preview :)** of what will evolve into the next NB interface: http://nb.csail.mit.edu:8080/?t=p23

I tried using the new interface, but it wouldn't link the comments with their location on the page. I would either click a comment box on the page and it wouldn't identify a comment, or I'd click a comment and it wouldn't identify the text it referenced. Did anyone else have this issue? — Update: after a third try, it worked. Not sure why.

No issues for me, sometimes it takes a few clicks to be linked to the comment

This example is pretty crazy. After reading the explanations I begin to understand the example and methods better, but for an introductory divide and conquer example, I feel the oil import question is a little better to start with.

It's great that this section is here, but many people were concerned about this when reading section 1.1. Perhaps you should make a note in that section that their concerns would be addressed in a future reading.

I agree. I think it would also be helpful to explicitly indicate that the capacity example will be supplemented by further approaches in section 1.2, so that as we read section 1.1 we can anticipate checking these assumptions through redundancy.

I disagree with these two comments and think it's fine the way it is. To us, it seems like it's a good idea to do that because there's been a break of a few day since we visited that material - however, to someone reading this book in physical form or as a whole entire PDF, it would be kind of redundant since they're only a page apart. Just flip back a page if you need a refresher.

I feel like you can never trust the estimate to make any important decisions. Then... I am confusing myself? Why do we make estimates anyways.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product 3.6×4.4×3.2. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so 3.6×4.4×3.2 is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5{\cdot}10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

Since we're working with N^a+/-b, how large is "embarrassingly" large? Or does that not apply in this case?

I think being off by more than an order or two of magnitude is when it begins to become embarrassingly far off, but I suppose this number gets larger as the numbers we're using get larger.

I agree with that comment. I think it's important to step back from a problem and ask yourself "does this answer make sense?" If you're getting a number that's physically impossible/unlikely (like a speed faster than the speed of light), I think it might be time to look at your assumptions/steps again to see if you can do better. This ties into the whole "gut" feeling that we talked about in class.

i've never done this, but i suppose it makes sense. i don't see how it's much different than just adding them again though.

displaying the list of numbers first would clarify this

Agreed - if you put 12, 15, and 18 in a list, and "rearrange" them to show the two different additions, the redundancy would be more clear.

I don't think I've ever used this method of checking my mental math, but it does make more sense than checking by doing it in the same order again. My way could allow me to make the same mistake twice in a row using the same thought process, similar to how you could type your e-mail incorrectly twice in a row as we discussed in class.

I've never heard of anybody doing this. Is this common? If I was going to check it, I would just add them again in the same order, or add them more slowly/carefully.l

It's better to do this way I think because that way you don't make the same easy mistake twice.

I agree. I often subtract some numbers from the sum and see if i can get the original numbers

I didn't know it was common either and did the same as you, adding them again in the same order or more carefully. It does make a lot more sense to go backwards. I can definitely remember some "mindless redundancy" mistakes way back when...

It's not clear to me what you're doing here

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product 3.6×4.4×3.2. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so 3.6×4.4×3.2 is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5{\cdot}10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

I've actually never thought of doing this, but I like the idea of it. I'll probably use this method in the future.

I agree, I have never thought of doing this before, but I think it could have really helped me in the past. I'll definitely remember this one the next time I have to check a sum.

Another good method for double checking sums/differences is the method of "casting out nines." I learned this as a child, and I found it useful because it didn't involve doing the problem again using the same method (adding the numbers forwards and backwards to me is the same method).

what is "casting out nines"?

The idea is that you add up the digits of what you're adding, and make every instance of "9" a "0,"then you add what is left (which is a single digit number). This answer should be the same as your original answer (assuming you add up the digits and "cast out" the 9s there as well).

For example, 193 + 324 = 517. I know this is correct. But if I didn't, I could use casting out nines to check. 1+9+3=4 (I made the 9 a 0). 3+2+4=9=0. So when we add up the digits of our answer, they should equal 4 (since 4+0=4). Alas, 5+1+7 = 13, and 1+3 =4. Therefore, our answer is most likely right. Of course, we could have made some weird error that this method wouldn't catch, but the method is usually correct.

The Wikipedia article on the method is decent–you should read it if you are unclear.

That's pretty useful. It might be beneficial to share it with everyone. It's at least a fun litle trick.

I've actually never checked my answer using a different method. This could be useful.

what exactly is "casting out nines"?

So I am having some troubles with this site now that I didn't have before; I cannot check all the comments. I want to see all the comments on an area, but can only see the largest one. Additionally, I am having global comment trouble. I cannot make or reply to global comments.

2. *Everything else.* What remains are the mantissas – the numbers in front of the power of ten. These moderately sized numbers contribute the product $3.6 \times 4.4 \times 3.2$. The mental multiplication is eased by collapsing mantissas into two numbers: 1 and 'few'. This number system is designed so that 'few' is halfway between 1 and 10; therefore, the only interesting multiplication fact is that $(\text{few})^2 = 10$. In other words, 'few' is approximately 3. In $3.6 \times 4.4 \times 3.2$, each factor is roughly a 'few', so $3.6 \times 4.4 \times 3.2$ is approximately $(few)^3$, which is 30: one power of 10 and one 'few'. However, this value is an underestimate because each factor in the product is slightly larger than 3. So instead of 30, I guess 50 (the true answer is 50.688). The mantissa's contribution of 50 combines with the eight powers of ten to give a capacity of $5 \cdot 10^9$ bits – in surprising agreement with the capacity figure on a box of CDROM's.

▶ *Find the examples of divide-and-conquer reasoning in this section.*

Divide-and-conquer reasoning appeared three times in this section:

1. spacing dissolved into capacity and area;
2. capacity dissolved into playing time, sampling rate, and sample size; and
3. numbers dissolved into mantissas and powers of ten.

These uses illustrate important maneuvers using the divide-and-conquer tool. Further practice with the tool comes in subsequent sections and in the problems. However, we have already used the tool enough to consider how to use it with finesse. So, the next two sections are theoretical, in a practical way.

1.2 Theory 1: Multiple estimates

After estimating the pit spacing, it is natural to wonder: How much can we trust the estimate? Did we make an embarrassingly large mistake? Making reliable estimates is the subject of this section.

In a familiar instance of searching for reliability, when we mentally add a list of numbers we often add the numbers first from top to bottom. For example: *12 plus 15 is 27; 27 plus 18 is 45.* Then, to check the result, we add the numbers in reverse: *18 plus 15 is 33; 33 plus 12 is 45.* When the

I agree with one of the global comments; how would you go about estimating something you know nothing about? Could you just say, "Well, I can't see the pits, so it's probably somewhat smaller than what I can make out visually... maybe a few nms?"

This goes with a lot of the problems on the pre-test; What are we supposed to do about estimation when we have no general knowledge on a subject?

I think a lot of us are adapting to this. You do need to start somewhere and I'm hoping to have a better feeling for sizes of things and more useful knowledge coming out of this class.

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

Comments on page 2

This is not always the case, especially if a computer is doing the calculations and spanning several orders of magnitude with small numbers first/larger laterr. We were learning about this in 2.086 (numerical computation) last week... something to do with the number of decimal places the program stores when doing calculations

What would be the possibility of both additions containing an error?

are there certain cases or pitfalls where this may not be true?

But doesn't redundancy only reduce careless mistakes? If we keep doing the math over and over, but our assumptions are wrong, our final estimation will still be off. Or does "intelligent redundancy" refer to coming back after we've gained new knowledge?

It does seem like there needs to be a distinction between redundancy to reduce arithmetic mistakes and redundancy to challenge assumptions through multiple approaches.

I had the same thought; its addressed in the next sentence

But I wouldn't call checking for arithmetic mistakes "mindless redundancy". For example, repeating the addition from the bottom up was a good way to check the addition, however it did not guard against wrongly assumed numbers to begin with. I realize the example is meant to be representative of the process as a whole, but I think it would still be useful to point out that there seem to be two types of helpful redundancy: reducing arithmetic errors through repetition, and checking assumptions through separate approaches.

I think there's a difference between mathematical redundancy (checking your work) and statistical redundancy (doing multiple tests to confirm results). I think this paragraph is trying to explain how to use statistical redundancy (different "directions" of adding) to reduce mathematical errors.

In the next few sentences, this concern is actually explained. It's natural for the mind to become accustomed to what's written and automatically assuming what we see is what we had in mind, and as a result, it's often difficult to catch one's mistakes without taking a step back and using 'fresh eyes'.

Down two paragraphs it explains intelligent redundancy as making the methods different from one another.

What level of redundancy is ideal?

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

Looks like we agree. (See my previous note.)

answered my comment/question above on pitfalls

I wonder if there are multiple methods of corrected errors in rough drafts by yourself without the time delayed by resting for fresh eyes, or asking others to read

I feel like this comparison might be a bit of a stretch. I can understand why your brain will assume the same answer to a calculation you did a few seconds earlier, but when you proofread a paper, it is usually very obvious when you make mistakes.

My experience is very different. Often I still find typos in papers I've spent years working on; and I usually find them when I have put down the paper for several months and then come back to it.

Agreed, for me I just have a problem spending too much time focused on paper/math problem. So the time away allows my mind to see the paper/math problem in a new light.

There are also those cases where you make one mistake so often that you're own brain has a hard time picking it out. This is especially common with poor spelling.

i agree, sometimes just the difference of looking at it on paper instead of a computer screen is enough.

This is a very common proofreading technique, as is reading drafts aloud. Sometimes if you talk your way through a problem aloud, careless mistakes become even more obvious.

Agreed. Both taking a break and reading drafts aloud can help. In 9.00, I learned that when you concentrate for too long on one thing, you are more likely to look over or make mistakes. However, switching your focus to another subject may not have the same effect because it activates a different part of your brain.

The switch to active imperative voice from the passive of the previous sentence was strange to me. "Putting the draft away for a week or asking someone else to read it, like adding numbers in reverse, provides a fresh look at the problem."

The analogy was helpful.

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

Or we can change the font and the size and we will be able to see the text in a "new" pictorial way, the brain will probably catch new characters.

Also a good strategy when debugging code. In 6.005 Professor Rinard suggested going surfing before coming back to the problem to have fresh eyes or a new perspective.

which robustness heuristic?

I am also confused as to which robustness heuristic being referred to,

The heuristic refers to mindful redundancy. I think it's just unclear- Ihad to read it a few times to understand that too.

You seem to be missing a verb. "This robustness heuristic was -used- at LIGO," otherwise it makes it sound like a physical object.

How does the system take advantage of this robustness heuristic?

I was curious as to what a gravitational wave was. The most succinct description I found was that bodies (stars, etc.) leak energy in the form of gravitational waves: Here's the Wikipedia entry: http://en.wikipedia.org/wiki/Gravitational_wave

I think NASA has a pretty good page on gravitational waves, describing them as "ripples in space-time." Here's the link: http://imagine.gsfc.nasa.gov/docs/features/topics/gwaves/gwaves

Wasn't sure if this was WA or DC so I looked it up. Its in Hanford, WA

not sure what the point of this is or what it mean.

what kind of acoustic noise? this seems like it would be everything...

I think that the acoustic noise does mean the background noise but also any other kind of noise fluctuation that would be site-specific and could be mistaken for a gravitational wave event, but not the kind of gravitational waves that they are looking for - they want ones that are global, not local.

This makes the class example of entering one's email twice seem insignificant in comparison. Having both does help clear up the idea, though.

After re-reading, I understand, but could be clearer

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

I think I understand this metaphor to mean that it's unlikely to make the same mistake over a variety of methods? However, it's a bit unclear, and I think the metaphor would benefit from being more explicit.

Oops...seems my comment didn't highlight the intended section - I was talking about the LIGO fact sheet example on page 2.

Yeah I agree. This example is a little cloudy.

The example made sense to me. Whether using different types of data (capacity, optics, hardware) or using the same type of data from discrete locations (Washington, Louisiana) the estimations are unlikely to agree unless they are reasonable. If the numbers are greatly different you know one of them must be inaccurate.

I think that it's just a little over kill to offer the LIGO example. There are already 2 'common' examples and adding a third that's less relatable to most people muddies the water.

Are there any tools we can use for approximating the level of redundancy that would be declared "intelligent?"

what was intelligent about this example? merely the fact that they were located far apart?

The intelligent use of redundancy in this case is the assumption that false gravitational wave events can be found (and ignored I suppose) because they will not occur in both, far-apart locations - however, the events that they want to find will occur at both locations

it seems that "intelligent" is a keyword in this section meaning "reasonable" or "intended by design". It is also repeated below as "intelligently redundant methods." However, I do not see "intelligent" being defined (although it might be somewhere), but it would help if it were defined.

I completely agree, but wouldn't it make sense to err on the side of being too redundant than not enough?

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

Is there anything we can do about the fact that the forces of nature are seemingly always against us? We assume that freak accidents are very very unlikely to happen, and yet they do...

This is a good point but if the system is constantly detecting gravitational waves then it should calculate a virtually accurate average measurement. This average would also make it easier to detect the microearthquakes because they would deviate from this number.

can the methods use completely disjoint knowledge and information though?

This seems like a very helpful and reliable process... however, for some problems it's hard enough just to find one way of getting an answer, let alone two or three different ways.

I agree; for the time being, we're having considerable trouble coming up with one trustworthy method, but perhaps over the course of 6.055 we'll be able to learn how to come up with a few of them relatively quickly.

Well, in the previous CD example we found, for instance, the storage capacity of a CD a few different ways. The entire problem may have one logical method of solving, but if we can tackle each constituent estimate from a few angles, they become more precise and robust.

But what if we do not have enough knowledge to try another way, and we can only run the numbers in one direction?

this reminds me of the diagram prof mahajan drew in class about gaining an understand about an area by exploring two unrelated examples

This is interesting...From lecture, I thought of divide and conquer as drawing from related sources that will influence your estimation. However using unrelated knowledge makes it easier for me to grasp. Initially, I was concerned about using divide and conquer because I assume I wouldn't know about every topic to make a estimate on it. ie. suppose I didn't know the population of the world...I figured that would hinder me in making a estimation of how many toothbrushs are used yearly (disclaimer: this is a random example)

How would you go about doing this, without knowing the population. I guess you could think about how much $$ a toothbrush company makes and the production/profit cost of a toothbrush, and how many companies there are...

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

It seems this is more of an "attack from multiple fronts" argument a than a "divide & conquer" one. I think the phrase divide and conquer gets misused quite a bit.

What about it is divide and conquer? It seems more like a "multiply and conquer" Since we are actually making more problems of the same length, rather than more smaller problems.

I think the "divide and conquer" similarity is that you want to divide it up different ways and see if you get the same answer.

Leave comments in the margins instead of overlaying over everyone else! It's very convenient to be able to pinpoint a commented section and click on it so the right comment bar centers in on the relevant comment...

Perhaps an improvement to NB that could be looked into would be prohibiting comments covering the same area or something. Although I think you should be able to comment on an entire paragraph and also a section of the paragraph with separate comments, this paragraph highlights just how bad it gets when overlap isn't prevented.

but then that defeats the purpose of being able to "pinpoint the area of confusion." but i agree, 75 people writing on one page is ridiculous.

Is there some other better reason for putting it in this chapter? This parenthetical seems unnecessary.

I agree with the comment above that this parenthetical insertion is unnecessary. In general, we would like to trust the author's judgement about the contents of a chapter.

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

This explanation of "divide and conquer" doesn't make sense to me. Hearing the name of that tactic makes me think that we are going to break up a problem into several pieces and conquer each small piece and then put them together to make an estimate. This idea of "divide and conquer" as a method of looking at several ways to solve one problem seems like its harder than just making one estimation because you have to find several ways to look at it.

I agree, I also thought this was a strange definition of divide and conquer. What seemed strange to me was that the CD example was just explained as a divide and conquer problem, but now we're being introduced to this heirarchy of divide and conquer problems. First, there is the CD problem, and finding an approximation for that is a divide and conquer problem. But now there is also this larger divide and conquer problem of obtaining a reliable approximation for the CD problem. Couldn't you infinitely come up with more divide and conquer problems which seek to refine the answer to another?

the idea of finding redundant methods is a _supplement_ to the divide and conquer tactic. it's a way to 'sanity check,' if you will.

Ok, so we have been using "divide and conquer" by splitting up each method of estimation into smaller problems, and then combining these smaller estimation problems in order to produce a "reliable" estimate. Are you also defining the use different methods of the same problem as a "dividing and conquering" method? It seems that we a dividing a checking here, not dividing in order to overcome the problem.

I think this is supposed to be a supplement to the divide and conquer method.

just because you split them up, does that guarantee that the error will be less?

It just means if there is error, then it is likely to be different in each place; if the answers don't match (for some reasonable definition of match), then you have at least one problem.

it is the same idea as adding numbers in different orders.

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

I like this explanation.

I think this confuses me more actually. What makes more sense to me is that you have divide & conquer to take a single estimation method and break it down into simpler subproblems. And then, outside of divide & conquer, you make your estimate more robust by repeating this exercise from a different approach. What the text implies (and what I don't understand) is how this robustness relates to divide & conquer, since it's not really dividing the problem so much as coming at it from a different approach.

I would have to agree with this. The explanation kind of muddies the concept of divide and conquer. I understood it to refer to the process of breaking a large problem into reasonable bits and then solving from there. The way this explains making the estimate more robust makes it seems like we are validating each component of the original estimate rather than the final solution. The examples, however take the problem as a whole, divide it differently and then conquer it differently.

I think this is a very valuable skill to be applied to engineering.

Would it be possible to include examples that are less technical but still get the concept and lesson across? It is sometimes difficult to follow technical-heavy examples such as this CDROM example (and the bandwidth 747 example in class).

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

Is this suggesting that we could have used any of the following methods (or any of the ones from the past reading) and we would come up with the same estimate or is it rather saying that in order to confirm/ disprove our past estimate that we should check it again with one of these? It was hard enough for me personally to think of way of looking at that problem- I'm really not sure I could look at it in several different ways unless prompted to (i.e. "Consider this now as an optics problem")

I agree, it would have been hard for me also to come up with another method of approximating this value. I wouldn't have been able to come up with either of the 2 following methods. Is there anything we can do if we aren't familiar enough with the subject material of the problem to come up with multiple methods?

I don't think that we should approach this as an 'exhaustive' list of methods - these are two methods that CAN be used, but don't necessarily have to be. Personally, i doubt I would have come up with either of them either, but that doesn't mean other methods that we may or may not come up with aren't applicable. My concern comes is at what point the error becomes too great, even with some redudancy precautions.

I would like to reword "intelligently" with "reasonable" in the document's sentence.

"optics-based"?

He uses 'based' later in the sentence.

I made a comment below for the analogous section of the second method. It is awkward to say "an optics method..." because too much information is being condensed into only 3 words. Instead, you might say: "A redundancy method that makes use of knowledge about optics is..." I proposed a similar phrase for the beginning of the second method. Another thing that my suggestion accomplishes is that it acknoledges that each method requires certain knowledge. Many students (including me) are complaining that these methods are very sophisticated and not very intuitive. By emphasizing that each method makes use of specific knowledge, you are telling the reader, "if you know A, you might do B as an option."

This method is pretty intimidating, I don't think I could do this.

me neither.

I don't know. These are equations that we learned in 8.02 and 3.091, the only part that really takes an intuitive leap is the measurement of the 0.5 rad.

two totals agree, as they do here, each is probably correct: The chance is low that both additions contain an error of exactly the same amount.

Redundancy, it seems, reduces errors. Mindless redundancy, however, offers little protection. As an example, if we repeatedly add the numbers from top to bottom, we are likely to repeat our mistakes from the first attempt. Similarly, reading your rough drafts several times usually means repeatedly overlooking the same spelling, grammar, or logic faults. Instead, put the draft in a drawer for a week, then look at it, or ask a colleague or friend – in both cases, use fresh eyes.

This robustness heuristic was in the Laser Interferometric Gravitational Observatory (LIGO), an extremely sensitive system to detect gravitational waves. It contains one detector in Washington and a second in Louisiana. The LIGO fact sheet explains the redundancy:

> Local phenomena such as micro-earthquakes, acoustic noise, and laser fluctuations can cause a disturbance at one site, simulating a gravitational wave event, but such disturbances are unlikely to happen simultaneously at widely separated sites.

Robustness, in short, comes from *intelligent* redundancy.

This principle helps us make reliable, robust estimates. Not only should we use several methods, we should make the methods different from one another; for example, make the methods use unrelated knowledge and information. This approach is another use of divide and conquer (which may explain why the approach belongs in this chapter): The hard problem of making a robust estimate becomes several simpler subproblems – one per estimation method.

So, to supplement the divide-and-conquer estimate for the pit spacing (Section 1.1), here are two intelligently redundant methods:

1. An optics method is based on turning over a CD to enjoy and explain the brilliant, shimmering colors. The colors are caused by how the pits diffract different wavelengths of light. (Diffraction is beautifully explained in Feynman's *QED* [12].) For a pristine example of diffraction, find a red-light laser pointer, the kind often used for presentations. When you shine it onto the back of a CD, you'll see several red dots on the wall. These dots are separated by the diffraction angle. This angle, we learn from optics, depends on the wavelength (or color): It is λ/D, where λ is the wavelength and D is the pit spacing. Since light

but if the pits are in a grid pattern, how do they look so smooth and connected rather than just a mix of a bunch of other colors

How interesting that the pits tend to refract into rainbow patterns so consistently.

Would it be possible for you to post the notes for the entire chapter so that we can look at the references?

This paragraph has two separate points: one, that you can use basic optics to get the pit spacing, and two, that the same principle is why CD's are pretty. The first is the main point, and would benefit from a bit more elaboration (perhaps a diagram). The second is an interesting side note, and should probably be a different paragraph.

A diagram of what this looks like would have been helpful here

I agree. Either a diagram or an in-class demonstration. Or both!

I also agree. It would be helpful if you provided a link to the demonstration if drawing it in the readings isn't possible.

I don't really understand how someone could think this up, and I agree that a demonstration would greatly help visualize this experiment.

so if were one were to use a green laser, the angle of the diffraction would be different, but would the pattern also change?

Random, but it would be really cool to test this using a green laser and a red laser, and see how the diffraction pattern shifts

I agree! It would also help us see how accurate our estimations are.

I don't understand how the D is automatically assumed to be the Pit spacing, what if the reflection pattern is cause by some different aspect of the CD such as the thickness, the coating, or the angle of the sides of the pits or anything along those lines.

rad

Not very understanding of these derivations. I'm not very acquainted with optics

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall. There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

Comments on page 3

But if your shining a red laser pointer, how can you get different colors out? The laser emits red at its wavelength and then refracts by the relation you mentioned. A red laser doesn't really have a spectrum of colors.

You are right, lasers should not emit a broad range of wavelengths, if any at all. Also, the sentence is not well written because it contains a false cause-effect relationship. That "lights contains a spectrum of colors" does not cause colors to diffract by their own angles.

Does this mean there would have been a different angle of diffraction if you had used green or blue instead of red in class today? I assume it would because it would change the lambda/D relationship, but it's hard for me to grasp an intuition for why that's true.

Different colors have different wavelengths (e.g. the visible spectrum), changing the values in the equation.

Are you talking about just shining white light onto the disc, or is this still referring to the red laser?

meaning that the colors look pretty? or they make interesting shapes?

I don't understand this leap.

Me too–is it because distinct red dots show up when a laser pointed is shined onto the CD? or is it because each color diffracts by its own angle? or some other reason?

I think the point is because there ARE so many colors, this implies that the colors are significant. Were the back of the CD uniform in color than we would not think to examine it.

I think poster #3's explanation is correct, but it would be great to have this more explicitly/thoroughly stated.

I don't know if brilliance is really the right word to use here either, unless it's defined explicitly to mean something like "the wide variation in color".

Agreed; I thought brilliance meant "brightness" here, and didn't realize it meant the variety of color until reading comments.

So simply the separation of colors is enough to conclude that the angles are significant?

also, what are we defining as significant? i probably wouldn't have come to the conclusion of 1 rad in difference.

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

Significant meaning that the definition of significant angles is 1 rad or that the angles are truly unique and then there was some math involved in order to arrive a the 1 rad?

wait how did we get to 1 rad?

I believe it has to do with powers of 10. A 'few' radians would be 3, or pi. This would correspond to an angle of 180deg, which wouldn't make sense for this reflection example. Thus, because we see such a large variation in colors, angles are on the order of 1 radian (60 deg).

I would also have preferred that this was more explicitly explained. I'm not sure I understand where 1 rad comes from..?

I agree that the rainbow effect means the diffraction angles are significant, however I don't follow the leap to 1 rad. I'm not entirely sure what is comparable to 1 rad.

Are the results of this test affected by the sunlight?

I just tried this with a CD and didn't see the "surrounding rainbow" (either that or I didn't know what to look for). If the point of this estimation is that anyone could do this, then I think either a diagram or a better explanation of what to look for would be helpful.

I've seen this experiment/calculation done in a freshman seminar and 6.007 - I feel like the presentation of the mathematics in those classes helped to understand what's being talked about. Can that be included in a future revision?

I had to read this part 3 times to figure out what it was trying to say...I think that images would really help me out here.

How do you compare this angle if they are reflected on the same surface?

I don't understand how they get that angle- perhaps a diagram?

an angle relative to what?

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

this section goes through approximately a class and a half of 8.02 in like a paragraph. it's very hard for me to follow and i'm still a little hazy.

So...the reflected light from the CD onto the wall made a CD-shaped bright spot, and then a rainbow (I'm assuming it stretched above or below the bright spot) that reached around 30 degrees, which is equivalent to 0.5 rad. Plugging this into theta yields 1/2 wavelength/pit space->pit space 2*wavelength. Given that the average value of visible light (which is what is seen in the rainbow on the wall) is about 0.5 micrometers, he plugged that into the equation pit space 2*0.5 to yield a spacing of around 1 micron

A figure or diagram would definitely help with the explanation of this.

Also, when solving problems, it is hard enough to come up with one method to solve them. Without having the proper knowledge to solve the problem in more than one way (like with the knowledge of diffraction in this case) how can we be redundant?

Your explanation was unclear, and this explains it much better.

I think it would be a lot easier to follow this paragraph if you've taken 8.03, however a diagram would definitely be appropriate here. I'm still not really sure what he means by "the rainbow appeared at an angle of roughly 30 degrees". In relation to what? Pictures Please!

I feel like some of these more complicated examples defeat their purpose in teaching about estimation since we're so caught up in trying to understand the physics a lot of us haven't learned that we sort of miss the point. Maybe using a simpler example that people can more easily conceptualize would be better for getting the estimation technique across?

I am also having trouble visualizing what is going on here. A diagram would definitely be useful in defining your angles.

maybe a modification to NB is a way in which figures and diagrams and references can be supported. Basically, if a person needs more clarification, he can click on a symbol besides the statement which will then open a diagram, figure, reference off to the side or in a pop up window.

I agree with this, although once you look past the complexities, it becomes apparent what point you are trying to make - you arrived at the same answer via a different method. While understanding the method more fully would make the example more easily repeatable for the reader, just a vague understanding gets the point across.

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

I also agree that pictures would be make things much easier to visualize. It also seems like some of the things we're estimating require knowledge that we probably haven't learned, and it'd be nice to develop some sort of intuitive feel of how to estimate these things with real-world comparisons rather than drawing equations from physics courses we've never taken.

agreed with all of the above. it's nice to space out all this info so its easier to think about. also pictures.

I agree with Anon#5 a thousand times.

I would agree with this comment. I found myself stuggling trying to remain focused on the point that you were trying to get across about redundancy, but kept getting distracted by understanding the details of the problem

I understand how these checks work, but they still rely on some assumptions about CDs that, unless you knew in advance, you'd have no way of verifying yet.

I don't get where the 2(lambda) came from

Is it saying D=2lambda because that would make sense

I have been feeling generally confused during these kinds of trains of thought, and I think I understand why. The specific figures make logical sense what I read them, but I doubt I could recall them on my own, especially the ones from the previous section about the pits in a CD. Can there be some elaboration on a systematic approach to a problem if you aren't sure of numbers to start with, or a way to figure out which numbers to use in certain circumstances? I realize this basically the heart of approximation and thus a pretty loaded question, but most of these examples have figures ready at hand, and it would be nice to see a problem that is approached with the same initial shock of knowing nothing (similar to the feeling I am faced with almost every problem on the pre-test).

As for the hardware method, I feel that it requires knowing/looking up information which makes the inquiry redundant. For example, if I know that the CD players use red light, I can easily find out the spacing of the pits.

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

I feel like in terms of understanding, this example works a lot better than the previous one since wavelength is something we can all conceptualize a bit better. (maybe it'd help to see this example first so by the time we get to the next one, we at least understand the approach?)

This wording is awkward because a lot of information is being compressed into only 3 words. When I saw "hardware" I thought "hammer and wood". Instead of abbreviating a concept you might say: "A different method that relies on knowledge about how a CD works as a computer hardware device is..."

Although I don't feel like this is common knowledge either, you have to remember that if we were to estimate something in real life, it would most likely be in a field related to what we are working on. In that case something like the wavelength of infrared radiation would be common knowledge

this sounds like a transcript of a lecture you're giving. since this is a book, i don't think "i seem to remember" is necessary.

Actually, I think he's emphasizing a situation where we might not have any tools or lecture notes and only vaguely remember some facts. The point of this class is to show us how we can mentally arrive at a guess on the spot. The way he is writing here seems to be literally simulating what might go on in our own heads, if we were doing the calculation the way he is trying to teach us to do.

I don't think the fact that the laser is near-infrared would be common knowledge, though it is entirely possible that many MIT students have a rough idea of the wavelength o red light - this brings up the question - what should be considered "common knowledge"?

As much knowledge as you can common! And the way to do that is to practice. One reason for not doing 'regular' grading in this class is that the base of knowledge each person has is so different, and my goal is not to grade people on their common knowledge but to encourage people to expand that knowledge.

yeah it does sound weird

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

initially I didn't understand what this section was talking about and I missed the "near" in the expression so I was very confused by the following sentences describing the terms. When I read it a second time however, it made a lot more sense and I think the hardware method is a really cool way to estimate the pit spacing. The confusion may have come from the fact that the boxes from people's previous comments make the text more difficult to follow because they are distracting.

I believe Blu-ray players work just like this except the laser is blue (smaller wavelength), thus it the disk can hold way more space

Correct: "While a standard DVD uses a 650 nanometer red laser, Blu-ray uses a shorter wavelength, a 405 nm blue-violet laser, and allows for almost ten times more data storage than a DVD." - http://en.wikipedia.org/wiki/Bluray

(Also note that the DVD has a shorter wavelength than the CD, as noted later in the paragraph).

Is there something preventing us from, therefore, making a system employing UV light and thus cramming even more onto a disk? What I mean is that I'm wondering whether it's a theoretical or purely technological problem.

I would not have know the details of how a CDROM drive reads data, but I like this section a lot otherwise. It is much simpler to follow than some of the more technical details in the previous section.

and my journal of numbers increases

I agree with the other comments about this section. Although I can understand how you got to your answer by reading this method, I lack the necessary "givens" to make these approximations.

How does the laser relay information back to the hardware, does it refract in varying directions and that signals the information?

Does the accuracy change as the ratio of the laser wavelength to the pit spacing changes?

How much smaller, or does that not matter? Is there some sort of equivalent to sampling and aliasing here, such that a much larger wavelength would not have worked?

Curious about this as well. Is it really wavelength or frequency at play here?

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

Is this considered lumping?

If the laser wavelength and the pit are similar sizes, Why doesn't the laser hit in between the pits often? I would think having a laser much smaller than the size of the pit would make it easier to be sure which pit the laser is looking into.

Why is there the uncertainty between 800- 1000 nm? I understand that it has to be greater than the red light wavelength but I don't get why the reading just says "800 nm or 1000 nm". I think it would read better as "between 800 to 1000 nm" as that denotes a range as opposed to 2 distinct values.

Why does the laser wavelength have to be shorter than the pit spacing? Also, you seem to be assuming that the laser and pit spacing are optimized for each other. Technically, couldn't the pit spacing be any distance less than the wavelength?

which was the design limiting factor? the lasers readily available for reading or the spacing of the pits?

1um isn't slightly larger than 1000nm, do you mean on the larger end of that range of 800-1000nm?

I think he meant the wavelength of laser is smaller than the size of the pit.

Right, but how does he make a jump from 1000nm (which is 1um) directly to 1um? If we assume the near IR scanner is at 1um, then this breaks apart. It seems very very rough

It comes from the assumption that the wavelength will actually be smaller than 1000nm while the pit spacing is slightly larger, but for the sake of simple math we'll just approximate them both to be 1um with the knowledge that the theory would actually work out in the end (that is, that even though we're approximating both to be the same length, we know they are actually slightly different).

It does seem very rough, and this method almost seems as if it's "tailored" to get the 1 um answer that all the other methods have arrived at. It's a little too convenient...,

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

The change in units here from nano to micro and back is a little confusing for someone who doesn't have a scale of them.

you should probably know that if you're reading this however

where did this number come from?

The 1.4 came from the previous section.

Isn't 1.4 significantly different than 1? Does this discount the first method?

Remember that we are judging accuracy by orders of magnitude. All guesses are based on values with around 1 significant figure. Based on significant figures, 1.4 is equivalent to 1.0 in this case. Maybe this fact should be noted since the beginning of the chapter and we should be reminded of it often.

Now having read the whole document, I feel like the addition part at the beginning wasn't quite the same, because the methods were too similar. I mean, you aren't simply suggesting here that we do our multiplication in a different order, you are suggesting entirely different methods. For the addition, shouldn't we have used entirely different methods? Like doing one in our head and counting out one on our fingers?

I sort of agree. There are several more drastically different methods for addition that can be used. One good one I like for adding N numbers is estimating an average value (call it A), and then keeping a running total of how much each number differs from that average (D). The sum is then N*A+D.

I think what this paragraph is getting at is that we come to similar answers with different methods, basically proving that our estimated answer is pretty much correct. The different methods are just a way to check your answer.

That's pretty cool.

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

After we reach these three results, assuming they differ, would averaging them out give us an even more accurate result? (My assumption would be that they aren't all under or over estimates)

i'm not sure either because one could also just say they all seem like 1micron and keep the 0.4 as our tolerance, right? or would it make more sense to say 1.1+/-0.3 or 1.2+/-0.2?

In chapter 7 we'll talk about how to combine estimates. But the basic idea is (usually) to use the geometric mean. So, if your two estimates are a and b, and they are both equally reliable, then your best estimate is sqrt(a*b). If one estimate is much more reliable than the other, then the best estimate will be closer to that value than to the other one (a weighted geometric mean).

How valuable can this be? How many types of problems really have enough information to be able to solve in a multitude of ways?

This phrase gets me to thinking. For this example, the estimates were all very similar in value. If the different estimates were not so similar, could the disparity between answers be factored in the uncertainty?

I don't think the answers should have to do with the uncertainty, but you may be able to adjust the uncertainty if you have underestimated it. For example, if I said that I had a 0.1 micrometer error after my 1.0 micrometer estimation, and then I got 1.4 micrometers the other way, I could probably assume I miscalculated the error.

How could such a consistent error have happened?

There could be some implicit assumption common to all three methods (lack of independence).

It didn't, I think. This example error is hypothetical.

I think this example is a little redundant. You have already mentioned in the intro to the section that the purpose of this exercise is to increase robustness and provide a "check" that your answer is probably close to correct. I think that it is obvious enough to the reader that being a factor of ten off consistently is unlikely on this scale.

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 μm to 0.7 μm – let's call it 0.5 μm – I estimate the pit spacing to be 1 μm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 μm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 μm (capacity), 1 μm (optics), and 1 μm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 μm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

Why is the error a factor of 10?

It means that the numbers are all in the same order of magnitude.

Does this mean that we can allow room for error as long as the results are consistent?

perhaps a discussion of systematic vs. random errors would be good here. In this hypothetical example, we just made a systematic error. In conclusion, redundancy cannot provide valid error-bars for systematic errors. Systematic errors might often be due to similarities between the redundant methodologies...

Is there a way to try to catch yourself, or check yourself, to ensure you haven't make some systematic error? Is it just a gut feeling?

The best way is to use methods that are very different. Then the systematic errors in each are likely to be different in type, and are not likely to point only in one direction.

So if the errors weren't coincidentally about the same, where would we go from there? Which estimates do we take as the more accurate if we didn't know the actual value? We wouldn't be able to keep doing this–given common knowledge, there aren't too many more ways to estimate this right?

I'm not sure knowing that CDs are read by near IR is common knowledge! And I have no idea where one would go next...maybe back to the basics? But this assumes you know the answer, so you'd have to find where the fudge factors went wrong

I think we'd at first have to look over all of our estimates and make sure we didn't make any foolish mistakes. And if they looked robust, look for a 3rd or 4th method of getting the estimate, perhaps by consulting with peers for ideas. For example, we could try another method that is not divide and conquer at all. Possibly some sort of dimensional analysis method.

Or one could evaluate all the assumptions made and see if one could have possibly have been made in error.

welcome to the dirty world of engineering, where you often need to know what your answer will be before actually obtaining it.

Yeah, what if the errors are not similar at all. Why give an example, and then say, "this example generally does not apply" without discussing an example that does?

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

With this sort of problem I would actually think these sort of errors might happen more often than this author seems to imply, most guesses are of a power of ten in one direction, and it is reasonable to assume that the guesses would be close to 1 order of magnitude off assuming the guesser is using some sort of logic, and so the chance that these errors might be similar should not be discounted. In the case of three I would guess it would be like flipping a coin three times, and saying either three heads or three tails means the errors line up, and in this case that would be 25%. In a case of 2 independent guesses, it would be 50/50. Obviously the magnitude that the guess is off would have a chance of being more than an order of magnitude so these % are a bit high, but I think it is reasonable for the guesses to all be wrong based on probability.

I would love for you to elaborate on this in class.

Does this mean we should repeat our calculation using multiple methods every time?

Would we then average the result and add an appropriate +/- tolerance, or make a gut judgment? Also, how far must a result vary from the others before we throw it out as bad data?

The second question is also unclear to me. Because the numbers here are so small, the different between 1.4 and 1 microns seems like a big enough difference to me to discount the first estimate.

I disagree that it is a big difference... they are the same order of magnitude which I think is more what we are after in this problem than an exact number.

I think we'd just do a quick mental average of the methods and forgo error bars, as error bars from averaging are negligable when compared to the errors from the approximations made.

I think we'd just do a quick mental average of the methods and forgo error bars, as error bars from averaging are negligable when compared to the errors from the approximations made.

It won't let me comment on the last page

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 μm to 0.7 μm – let's call it 0.5 μm – I estimate the pit spacing to be 1 μm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 μm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 μm (capacity), 1 μm (optics), and 1 μm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 μm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

Thats a funny expression. Where does that come from? Sounds like a British expression.

Heh, yeah it is a weird one. Never heard that before. I do like though how the text repeats the main idea multiple times, without being annoyingly 'redundant'.

the phrase is "the proof is in the pudding".... its just adapted here i guess

"the proof is in the pudding" reminds me of the slogan of a Yale all-female a capella group I met last semester, hahaha.

How long? I estimated 10km.

I agree – The rings are 1um apart and they cover about 3cm radially, so there should be about 3e4 rings. Average radius of a ring is about 4cm, so average circumference is about 20cm. That give us about 6e3 meters, or close to 10km total.

Where are these numbers from?

The 1um spacing is from the previous reading, and the 3cm radial coverage and average radius are from the ruler and CDs in my desk drawer.

I was the first poster–I used the lattice approach and thought of it as a 10 cm square of pits 1 um apart. One row across would be 10 cm = 10^-1 m long, and there are 10 cm/1um = (10^-1 m)/(10^-6) m = 10^5 rows, so the total length = 10^5 rows * 10^-1 m/row = 10^4 m = 10 km. The first hit in Google for length of track on cd says 6 km, so we're in the right order of magnitude.

contains a spectrum of colors, each color diffracts by its own angle. Tilting the disc changes the mix of spots – of colors – that reach your eye, creating the shimmering colors.

Their brilliance hints that the diffraction angles are significant – meaning that they are comparable to 1 rad. To estimate the angle more precisely, and lacking a laser pointer, I took a CD to a sunny spot and noted what appeared on the nearest wall: There was a sunny circle, the reflected image of the CD, surrounded by a diffracted rainbow. Relative to the reflected image, the rainbow appeared at an angle of roughly 30° or 0.5 rad. This data along with the diffraction relation $\theta \sim \lambda/D$ implies that the pit spacing is approximately 2λ. Since visible-light wavelengths range from 0.35 µm to 0.7 µm – let's call it 0.5 µm – I estimate the pit spacing to be 1 µm.

2. A hardware method is based on how a CD player or a CDROM drive reads data. It scans the disc with a tiny laser that emits – I seem to remember – near-infrared radiation. The *infrared* means that the radiation's wavelength is longer than the wavelength of red light; the *near* indicates that its wavelength is close to the wavelength of red light. Therefore, *near infrared* means that the wavelength is only slightly longer than the wavelength of red light. For the laser to read the pits, its wavelength should be smaller than the pit spacing or size. Since red light has a wavelength of roughly 700 nm, I'll guess that the laser has a wavelength of 800 nm or 1000 nm and that the pit spacing is slightly larger – 1 µm. (The actual wavelength is 780 nm.)

Three significantly different methods give comparable estimates: 1.4 µm (capacity), 1 µm (optics), and 1 µm (hardware). Therefore, we have probably not committed a blunder in any method. To make that argument concrete, imagine that the true spacing is 0.1 µm. Then three *independent* methods all contain an error of a factor of 10 – and each time producing an overestimate. Such a coincidence is not common. Although any method can contain errors – the world is infinite but our abilities are finite – the errors would not often agree in sign (being an over- or underestimate) and magnitude.

The lesson – that intelligent redundancy produces robustness – seems plausible now, I hope. But the proof of the pudding is in the eating: What is the true pit spacing? It depends whether you mean the radial or the transverse spacing. The data pits lie on a tremendously long spiral

Given that these are different (I didn't know that) couldn't we have taken it into account when we did the first estimation by making our square have unequal numbers of pits between column and row?

Yes, you could have. But doing so would make the problem harder and defeat the purpose of getting a simple "guess" at the pit-spacing. I think so long as we get something within the order of magnitude, it is "good enough". Introducing complications at the beginning will slow us down. This isn't to say that you shouldn't consider the fact that they are spaced differently transversely versus longitudinally, but that you should arrive at a simple answer first, then modify it using additional knowledge. It's like making a sculpture: you make big cuts to get the shape approximately right, then you start using smaller tools to carve out and edit the chunk into a more detailed/accurate representation.

I really like your sculpting analogy!

I agree that for a good guess, it is not necessary to actually calculate row and column spacing of the pits. However, it would be good to be aware of this difference since the beginning. In this case we were very fortunate that the spacings were very close. However, if they were far apart (eg. factor of 100), then we would be looking for a single number that has no physical significance at all.

I understand what this looks like, but a picture or figure would be a cool addition to the words.

What does this spiral look like? I'd understood it as concentric circles and not a single line curled into a spiral.

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition

The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy

Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

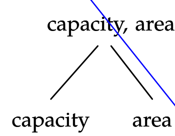

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

Why are there different spacings for the different axes?

I feel like compared to the previous estimations, this answer is either very easy/ brief or not explained. Should I know what is meant by 'rings' when we've only been talking about 'pits'?

he's making the assumption that you can visualize the cd: "the data pits lie on a tremendously long spiral track..." the spiral forms "rings" around the cd, not individual rings, but still rings

These two values measure the pit spacing in two different ways, "radially" and "linearly." Was there a specific way that the different estimates done before tried to measure the distance? Like maybe the estimates are more accurate than explained because they are measuring the distance one way versus another?

I think that these measurements are actually given (by manufacturers) data points that we are comparing 'our' estimates to.

ummm... ok. is this helpful to our understanding?

While I really enjoy the casual tone of this text (it makes it much more enjoyable to read), I think this sort of statement is a little over the line, and not as useful as a conclusion that pointed out how some of our calculations overestimated, and others underestimated (or some other conclusion). That's just my humble opinion, though. - Edit: I'm not sure if the tag came through, but this was referring to "We made a tasty pudding!"

We've defined intelligent redundancy, but what is unintelligent redundancy? Is it just when we repeat what we've done before or are there ways to be redundant without repeating your previous methods?

Great line.

I don't know if it's just me or if NB is broken, but none of the links on page 10 are working for me. The other 3 pages are working fine, just not this one.

I'm having trouble. Does anyone know an example using probability?

typo....also, why are no note boxes showing up on this page?

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

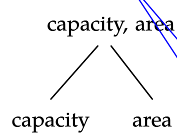

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

I wasn't sure how to analyze the robust addition practice in terms of probability models, could someone please clarify/answer this?

But, I do agree that robust addition is a good example of intelligent redundancy, whether it involves going from top to bottom, bottom to top, or grouping similar terms (like 3s and 7s, 4s and 6s, etc.).

I think it has to do with the probability of missing the same thing twice...if you increase the number of times and directions you look at the same thing, the likelihood of missing the same value decreases. Also, it gives more data sets to average over and find mistakes. I don't know much about probability models though.

In the previous version of the chapter, there was a section here on making probability models to determine the accuracy of an estimate. But I recently moved it to the new chapter on probabilistic reasoning. The question about probability models belongs there too. Thanks for pointing out that it is out of place.

Typo. Is it helpful for us to point these out?

Yes, thank you. The reading memos are a form redundancy. I've been staring at all these sentences for so long that I do not see the mistakes (whether typographical, mathematical, or conceptual). But all your fresh eyes find so problems so quickly. Sigh!

The probability of adding wrong going down is at least a little different than going up.

P(down) intersect P(down) = P(down)

P(down) intersect P(up) < P(down)

Such as coin tosses and picking marbles out of a jar? If you need enough repetitions, then yes I think it's intelligently redundant – right? But is it intelligently redundant with only a small number of repetitions?

what do you mean here by simple probability models?

It is smart because repeating an arithmetic mistake with different numbers is way more improbable than repeating the same mistake with the same numbers.

track whose 'rings' lie 1.6 µm apart. Along the track, the pits lie 0.9 µm apart. So, the spacing is between 0.9 and 1.6 µm; if you want just one value, let it be the midpoint, 1.3 µm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

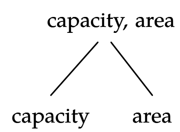

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

I feel like adding backwards and forwards and in whatever order isn't really intelligent redundancy–you're not changing methods or gaining any new knowledge...

True, but it is a way of trying to make sure you don't accidentally skip the same number twice

I also feel like it is very different. to add 12, 15 and 18 forwards requires you to (looking just at ones column) add 2+5, then 7+8, whereas adding in reverse requires 8+5, then 3+2. These are all very different additions. You might get the same result, but how you arrive at it is very different in the two cases, which is the point of intelligent redundancy.

I do not know what you mean by probabilistic methods here. Maybe it means, if you have a 10% chance of messing up an addition problem, you have a 1% chance of messing it up twice and an even smaller chance of messing it up to get the same number twice. ??

When proofreading my papers, I usually don't have time to set them aside for a few weeks and, when the paper is 20 pages long, it can be hard to find someone else to read it for you. So i usually read it in my head once, and then again out loud. When I read them out loud I sometimes notice things I didn't originally. "Hey...that sounds weird..."

I think this is one way, I've used intelligent redundancy before. When first learning to play the violin, intonation was something that my instructor worked on with me. Say I was learning where to place my finger to play an E in the first position on the A string. I would play it once, compare to the E the instructor played on the piano, compare it to the sound of the open E string, and use a calibrated tuner to help me adjust. In the end I would be able to play that E note in tune.

I still don't think I'm totally clear on "intelligent redundancy"–is it just making more estimations using different methods? I feel like this isn't really possible most cases becuase we don't always know so much about what we're estimating.

Yes, I think that you're on the right track. He is saying that by using intelligent redundancy, we can check our cross-check our answers to make sure they are relatively consistent. I agree with you that this isn't really possible in most cases given what we know yet, but perhaps we will have a better intuition at the end of the class.

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

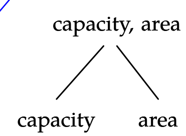

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

I thought the example about entering your email address twice when signing up for websites was really effective in class. It might be worth to add that as an example at the beginning of the section in order to provide a simple easy-to-understand example.

measure twice cut once?

Peer editing, grading itself..

In 6.02, we were taught that electronics read voltages as 1's and 0's (binary code). However, in real life, changing the output voltage from zero to one is not instant, it requires many samples because the transition is slow. The redundancy of the samples facilitate the voltage to reach its value

So basically we can try to increase the robustness of our estimates by going top down and then bottom up. Nice

track whose ‘rings’ lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition

The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy

Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate’s structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

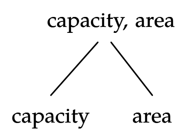

As an example, let’s construct the tree representing the elaborate divide-and-conquer estimate for a CDROM’s pit spacing (Section 1.1). The tree’s root is ‘capacity, area’, a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree’s root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the ‘area’ node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

Are nodes and leaves the same thing?

A Node refers to each individual "quantity" that we are finding. A leaf however is a node that does not sprout a new branch. In this tree, ’capacity’ is a node while ’sample rate’ is a leaf.

lateral movement represents "intelligent" redundancy while longitudinal movement represents divide-and-conquer. I like the trees because they are a more visual way of showing/remembering where numbers came from...though I agree that it would be nice if the longitudinal movements were labeled with equations or something of that sort. I can see these trees becoming quite complicated for other examples though

I thought this section was written pretty clearly. The diagrams make it easy to grasp the concept. Perhaps the wording could be improved. Like someone else said earlier, the first 2 paragraphs are awkwardly worded. Other than that, I have no real comments to add! I hope that’s not a bad thing.

I definitely like the visualization of the method. It makes the thought process more clear for people who like pictures.

I don’t think this actually helps in estimation. I feel like the difficulty in a lot of what we’re doing is figuring out exactly how we’re going to estimate something, and I feel like these trees, although new (so it may take some getting used to), maybe just be a wasted way of organizing something that we should be able to handle at any mildly high level of estimation.

I disagree. While many approximations that we’ve done have been simple enough that this is not necessary, I can imagine more complex ones for which this would be useful. However, much of its use is in looking back on the work you’ve done when it’s no longer fresh in your mind. You can follow your own logic (or explain it to others) very easily with a diagram of this type. It also makes errors easier to find and fix.

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

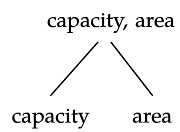

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

The teaching style of this class is so different than in other classes. I think it is interesting that all of the lessons from the reading all use the same example of a CD-ROM. Would it be more beneficial if the text had a different example or used more than one (maybe its just that I don't find CD-ROMs very interesting)?

I agree. But this is only based on the first three readings. I don't expect the readings to continue to be entirely about CD's. Look at example 1.4, which looks much more interesting to me, it has no CD's.

while i realize that, technically, there haven't actually been that many book sections on CDs, CDs have been the major focus for about a week's worth of class. i'm not very comfortable with 8.02, waves, digital information sizes, and electrical engineering, so despite sanjoy's best efforts, i still find CDs very difficult to understand.

i feel like, thus far (at least in the readings), i have been struggling more to understand CDs than to understand approximation methods.

This comment and many others along the same lines – all of which I agree with – have pretty much convinced me to significantly rewrite these first few sections. The main idea is divide and conquer, and it's getting lost in the details in my current organization.

I think the first example will be, as suggested on one of the paper end-of-lecture A5 sheets, estimating the number of seconds in a year. And use a tree from the beginning.

Once divide and conquer is introduced, then maybe do something like the capacity of a CD or the bandwidth of a 747 filled with CDs.

Why is looking at the box a viable option? Can't be deceiving? Also, how does that really teach us anything about the item/measurement? Shouldn't it be an option we use to make other measurements?

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

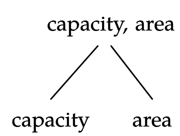

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

It would be really nice to see an example that hadn't already been done in addition, or maybe before, the cd one that we are familiar with. It could be helpful to see both how old ways of thinking about it translate and how the tree can help you think differently.

I agree to a certain extent... the CD example in my opinion is a little too technical for someone to be able to think up off the top of their head... but its nice reiterating a process that was already done for sheer sake of understanding the tree approach. Because of this, it allows the reader to be able to relate to something previously done.

So does the tree just serve to help you break down equations? Usually to make the tree what I looking for and then think of equations I need to find it and place the variables in the equations in branches below.

track whose ‘rings’ lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate’s structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

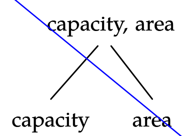

As an example, let’s construct the tree representing the elaborate divide-and-conquer estimate for a CDROM’s pit spacing (Section 1.1). The tree’s root is ‘capacity, area’, a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree’s root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the ‘area’ node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

Are you going to include solutions to these problems (for readers using the book for self-study, for example)?

Read this section for Wednesday’s lecture; the memo due Tuesday at 10pm. I’ll have the links from the 6.055 website point to the new NB user interface (the old UI seems to be a bit fragile right now). Let me and Sacha know if you have problems with NB.

Is there a way to get back to the "File" and "Nav" tabs on the left frame and the "Home" tab on the right frame after clicking to see our own comments besides clicking "Back" on the browser? I tried closing the "Controls" and "Notes", which resulted in just white empty frames.

I’ve forwarded that question to Sacha, who hopefully can say how to do it or will implement an answer!

I don’t feel like this section adds much new information. The tree structure seems helpful but doesn’t need much explanation. The bigger thing that I think will trip me and others up is not having the numbers at my fingertips with which to do these estimates. I feel like a list of numbers to memorize might be a good start to this class. Even if it’s boring, it would help give us a foundation of numbers to work from.

Meh...I disagree to a certain extent. If you just memorize things, you might start thinking too rigidly. It’s more useful if you just play around with things and estimate based on your intuition. Eventually, this playing around will allow you to naturally memorize certain values/constants w/o the pain of memorizing them for memorization’s sake.

estimation pudding? I have no idea what this is referring to...

It’s a metaphor meant to capture your attention and make the reading more entertaining. Or just make you hungry.

It’s a metaphor meant to capture your attention and make the reading more entertaining. Or just make you hungry.

The dash here seems a bit awkward, though maybe I’m reading too much into it.

It kinda makes sense to me, like a mental break

track whose 'rings' lie 1.6 µm apart. Along the track, the pits lie 0.9 µm apart. So, the spacing is between 0.9 and 1.6 µm; if you want just one value, let it be the midpoint, 1.3 µm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

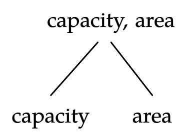

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

This whole paragraph reads a little awkwardly for me. The present tense is a little odd, though I do understand the intent. I feel like the current content could be shortened into a couple sentences and the rest of the paragraph fleshed out a tad.

I agree, I think the following argument reads well without this paragraph and wouldn't lose any meaning if it weren't there.

I'm actually just confused about the purpose of this paragraph. Is it just explaining how to think about divide-and-conquer?

I agree - I think the whole paragraph could be scrapped and there would be little impact lost.

I think this part is fine. It is not that informative, but a textbook that is information dense is too boring. This adds some spice to the whole thing.

I agree, the paragraph doesn't really explain much. It gives a bit of a transition into a better way to represent divide and conquer I believe but it seems a bit unnecessary.

These two beginning paragraphs are unnecessary and awkward. This whole part could be replaced with something like: "Estimation problems can be difficult and complicated to solve, and often have a hierarchical structure. For such problems, a tree structure can be a useful representation."

ˆagreed.

I disagree.. I find in these readings the conversational tone keeps it interesting. I personally enjoyed the personality in the intro.

This paragraph transition is also a bit weird sounding - it definitely makes sense, but I think it could flow a bit better to make the point that the best method of truly representing the estimate is a tree.

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

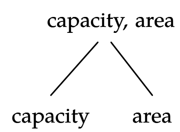

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

What does this mean? It reads very abruptly.

I think what it's saying is that in previous methods, we thought of the estimation methodology as very sequential and simply doing one step after another. What this Tree representation methodology is trying to improve is organizing the different components of our estimate so that it is easier to understand, rather than doing step by step where you could easily get lost.

This was also very confusing for me...I think use of more proper nouns would make the introductory sentence of this paragraph less ambiguous.

Linear and sequential also seem redundant (is there a difference in this context?)

This paragraph helps to introduce the trees and why they might be so useful. I think its a good way to introduce it.

What is the motivation for finding a fitting structure for this estimation process? I might suggest spending a moment saying why we want this whole thing to be visualized in a structure at all before diving into the tree details.

I totally agree with this. A tree diagram would definitely be a more efficient way of displaying the problem. When reading the previous section, I could not remember the process the author went through to get the answer. I will be using this for study sheets now

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

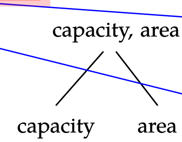

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

How does this fact bear on the actual procedure used in estimating? Should we write the tree as we work, or produce it afterward?

I agree - when should the tree be created? Is there anything in particular we should be looking for when producing the tree?

I think it's more the idea of the process. Going along with the recipe analogy, some cooking plans you have ahead of time. People might know for example that pancakes are going to go well with syrup before they are made. Other times you might stumble upon toppings that go well that you may not have expected - cooked bananas for example. You add these to the recipe as you go. Still other times, you might just be trying out new stuff. Then when everything works out, it's good to document it so you can do it again. I think it's flexible like that. Now you just gotta learn how to flip pancakes.

I feel like it depends on how you think. I would probably go about drawing the Tree as I worked through the problem. I also feel like how you draw the tree depends on your thinking style. As much as I don't want to, I can compare it to breadth-first and depth-first search. If I were to draw a Tree, I would do it depth first, working completely through an example and then moving on to the next.

I use a tree as I estimate, but most often I do it in my head. Maybe I do it in my head because I have lots of practice. So, my suggestion is to elaborate the tree on paper as you work. With practice you'll end up doing lots of the elaboration in your head.

Stating that the tree is the "most compact representation" says that there are other ways of representation. Do those other ways refer to linear and sequential or other representations for hierarchical information?

Are there similar constructs for the other methods we will develop?

For several of them – lumping, springs, easy [extreme] cases – a very useful visual representation will be graphs. I have lots of examples and will share those. I'll also keep looking for visual restructurings that are as significant as the tree representation is for divide and conquer.

shouldn't the 'root' be what we're trying to find?

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

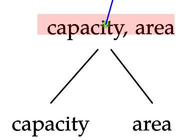

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

Why is capacity, area the root instead of pit spacing? What part of a generic problem will typically constitute the root?

I agree that capacity, area is not a particularly good name for the top of the tree. It adds no information since the two branches are themselves "capacity" and "area". "Pit spacing" gives some explanation of why you might want to find capacity and area, so you don't lose track of what you're trying to do.

Would we normally construct a tree after performing our estimations, or should we construct a tree as we work through a problem?

From what it seems, a tree method is just a way to organize our calculations so theyre easier to visualize. But we'd still have to come up with our whole plan beforehand, otherwise we'd end up drawing multiple trees and then putting them together (which should be ok too, right?).

For me, the tree seems like a helpful place to start. I am happy that this paragraph goes back and addresses a previous example so we see where we could have broken it down.

My guess is it depends on style. If you can see easily what needs to be done, then there is no need to draw a tree. But if you're lost, the tree is probably a very helpful tool to visualize where you could go and what you might be missing.

I think the tree just provides some visual structure with which to organize our divide-and-conquer thought process.

For me, this is something useful for me to do as I go along, even if I do have to revise it. I really learn from diagrams and need the clearer organization... for me this section is more helpful/clearer than the previous ones.

This goes back to the very first paragraph of section 1.1 when he divided the problem into the a) The problem of estimating the area of a CD and b) the problem of estimating the capacity.

I feel like the tree is just a way to visualize the problem in a different way. It all depends on what you're comfortable with and what works best for you

can this same idea be utilized for all divide and conquer problems? and how can you determine if this is a effective strategy?

track whose 'rings' lie 1.6 µm apart. Along the track, the pits lie 0.9 µm apart. So, the spacing is between 0.9 and 1.6 µm; if you want just one value, let it be the midpoint, 1.3 µm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

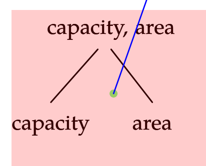

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

Is the main advantage of the tree representation that the pieces of the divide-and-conquer estimations are easier to visualize?

I really like the use of the tree method. I think beginning the divide and conquer explanation with this structure is a great idea and should be put sooner in the paper.

i think so and it's easier to break a complicated problem into multiple parts, which are often easier to analyze

This seems like a more reasonable way of explaining the pro's of redundancy, though seems like it might get tedious if there are many things to estimate.

Shouldn't the root be pit spacing? The two instances of capacity and area seem redundant.

I think the text might be trying to distinguish between other approaches such as "optical" and "hardware" as discussed in sec 1.2, so that if we added a "pit spacing" root node, "capacity,area" could be one if its child nodes.

I agree, but it still seems that "capacity, area" is overly redundant and unoriginal. Perhaps something like "Storage" or the like, which implies both capacity and area, would be better.

It might help if there were some way to flag here at this point that this tree will be attached to one with "pit spacing" as its root (as happens later in the section). For the label "capacity, area" makes more sense in the context of the larger tree than it does on this standalone tree.

agreed

maybe the tree diagram is only used to describe your problem, because now we can take this information and solve other things related to CDs instead of just being limited to finding the pit spacing

track whose ‘rings’ lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate’s structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

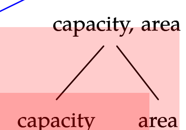

As an example, let’s construct the tree representing the elaborate divide-and-conquer estimate for a CDROM’s pit spacing (Section 1.1). The tree’s root is ‘capacity, area’, a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree’s root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the ‘area’ node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

So this tree is hanging from the ceiling?

this is standard programming language, and a standard diagram

Typical trees used for analysis typically start from the top down with the parent node at the very top. Another way to do a tress is from left to right. Either way, the point is to set up a hierarchical structure so you can relate certain "leafs" together.

It’s referring to a graphical and mathematical, as opposed to biological, tree.

And even in biology, pedigree charts are arranged in very much the same fashion, with the parents forming the root node at the top and children forming nodes below.

This reading clears up a lot of the confusion I had when I was reading the later notes. I guess I was in a hurry catching up to the class that I read some out of order, and that confused me a lot. Now it makes a lot of sense.

Is it the specific example or are we going to use the same names for the root and branches.

i think it is the specific example in this case, however it does appear in a lot of problems i think it’s a somewhat general approach

Isn’t sprouting capacity, area into capacity and area a bit redundant? Shouldn’t capacity and area come from some higher category describing the end goal?

That’s a good point - by using two tags that are the categories, that step seems pointless. However, I think in real world use, we could use the ’root’ for tags beyond just what the resulting branches are. ’tens’ and ’ones’ don’t have to trace back to ’tens, ones’. I suppose you could call the root ’cd guess stuffs’ but that’s obviously not very useful compared to ’capacity, area’.

As an additional note, when you get to the end of the lesson, you can see that this practical approach is useful because you may wish to expand your trees later with supersets.

I think the "capacity, area" tag is supposed to help you identify the method used, and the tag doesn’t necessarily have to be the combination of its two branches. But it should be more descriptive than "cd guess stuffs" in order to distinguish itself from other estimation methods.

track whose 'rings' lie 1.6 µm apart. Along the track, the pits lie 0.9 µm apart. So, the spacing is between 0.9 and 1.6 µm; if you want just one value, let it be the midpoint, 1.3 µm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

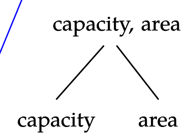

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

How can we decide which operations are trivial and which are not? Because here we are not actually estimating the area, we are estimating the length. I think we should be explicit here that we estimated length (and made it into a square) since that could be a considerable source of error.

it's trivial if you can arrive at an answer mentally. it was originally described as a grid, so that seems to fit the 'area' category fine.

It seems that the leaf can remain a leaf if you do not need to think of new estimations to produce the answer, in this case the only thing you need to estimate is the length and simplify square-ness. If you were to make this a branch, it would only be connected to area –> square simplification –> length; which is not very useful and really just 1 estimation. The other branches from the rest of the section are actually separate estimations. You need the playing time, sampling rate, and sample size (All estimations alone) to calculate the audio content.

I thought it was interesting that the first comment brought up error. Would it be possible to use how many nodes the tree has (and thus how many separate estimations must be done) to roughly estimate the error/exponent associated with our answer?

How? Each problem requires a different number of estimations, and doesn't this bring up the original question of whether one should estimate or leave the node as is?

I'm glad you asked! In the chapter on "Probabilistic reasoning", we will discuss how to estimate the error in one's estimate - and by extension how to choose which leaf nodes are contributing most to the error. With that information, one can plan where to put more effort into refining the calculation.

track whose 'rings' lie 1.6 µm apart. Along the track, the pits lie 0.9 µm apart. So, the spacing is between 0.9 and 1.6 µm; if you want just one value, let it be the midpoint, 1.3 µm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

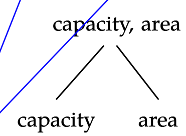

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

why do you think it's easier for us to estimate some values, such as area, than others, such as force?

> I think it's the amount of physical experience, plus misconceptions from ordinary language. For example, the ordinary use of 'force' overlaps with the physics concepts of energy, force, and power. So to start with it's hard to have a good estimate of a physics force.
>
> Second, we use area a lot more than we use force. So the fundamental units of area are much more familiar than the units of force. For example, I instantly can see 1 mˆ2; but 1 N? That I have to think about.
>
> Then I remember a useful rule of thumb: 1 N is the weight of one (small) apple. Perhaps the kind of the apple of legend that fell on Newton's head while he was sitting under the tree contemplating gravitation.

when is it a good plan to utilize intelligent redundant? when is it necessary?

How is this redundancy, if (1) solves it for you right off the bat? It seems beyond redundant, if that makes sense.

I am curious about this, too. Are we assuming that we can't trust the manufacturer? Or are we looking at the box to check our answer? If the box is already available to us, then why are we trying to estimate the capacity using other methods?

It seems like this is an example and you happen to have the manufacturer's information, but if you had another source instead it is good practice to keep this structure even if you are confident in one of the methods used.

I'm fairly certain that the box is implied to be someone's attempt at the answer (possibly a friend's guess). The redundancy is in the estimation.

Now in this tree context I understand this a lot more. and even the idea of intelligent redundancy can be applied in the tree method

I like how it is made clear that a tree can be used to either split some component of a problem into parts, or split it into different ways of calculating the same thing. (intelligent redundancy)

Yeah the tree makes it a lot easier for me to wrap my head around the CD example.

track whose 'rings' lie 1.6 μm apart. Along the track, the pits lie 0.9 μm apart. So, the spacing is between 0.9 and 1.6 μm; if you want just one value, let it be the midpoint, 1.3 μm. We made a tasty pudding!

Problem 1.3 Robust addition
The text offered addition as an example of intelligent redundancy: We often verify an addition by by redoing the sum from bottom to top. Analyze this practice using simple probability models. Is it indeed an example of intelligent redundancy?

Problem 1.4 Intelligent redundancy
Think of and describe a few real-life examples of intelligent redundancy.

1.3 Theory 2: Tree representations

Tasty though the estimation pudding may be, its recipe is long and detailed. It is hard to follow – even for its author. Although I wrote the analysis, I cannot quickly recall all its pieces; rather, I must remind myself of the pieces by looking over the text. As I do, I am reminded that sentences, paragraphs, and pages do not compactly represent a divide-and-conquer estimate.

Linear, sequential information does not match the estimate's structure. Its structure is hierarchical – with answers constructed from solving smaller problems, which might be constructed from even solving still smaller problems – and its most compact representation is as a tree.

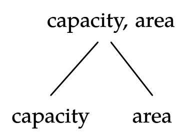

As an example, let's construct the tree representing the elaborate divide-and-conquer estimate for a CDROM's pit spacing (Section 1.1). The tree's root is 'capacity, area', a two-word tag reminding us of the method underlying the estimate. The estimate dissolves into finding two quantities – the capacity and area – so the tree's root sprouts two branches.

Of the two new leaves, the area is easy to estimate without explicitly subdividing into smaller problems, so the 'area' node remains a leaf. To estimate the capacity – rather, to estimate the capacity reliably – we used intelligent redundancy: (1) looking on a CDROM box; and (2) estimating how many bits are required to represent the music that fits on an audio

I understand the concept of intelligent redundancy, but I have not yet heard a good explanation of what to do when the guesses differ greatly. In all of these examples all of the numbers work out, but what should we do when it doesn't work out so nicely when we are solving problems.

You should probably recheck both and then find an alternative way.

CD. The second method subdivided into three estimates: for the playing time, sample rate, and sample size. Accordingly, the 'capacity' node sprouts new branches – and a new connector:

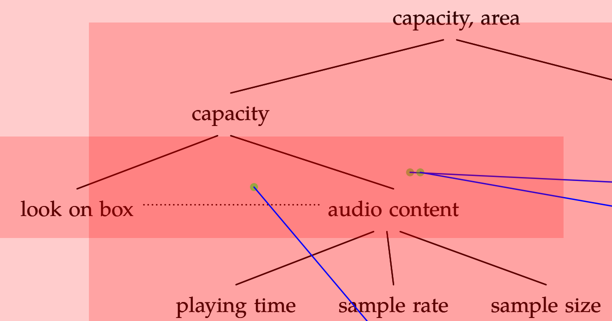

The dotted horizontal line indicates that its endpoints redundantly evaluate their common parent (see Section 1.2). Just as a crossbar strengthens a structure, the crossing line indicates the extra reliability of an estimate based on redundant methods.

The next step in representing the estimate is to include estimates at the five leaves:

1. capacity on a box of CDROM's: 700 MB;
2. playing time: roughly one hour;
3. sampling rate: 44 kHz;
4. sample size: 32 bits;
5. area: $(10\,\text{cm})^2$.

Here is the quantified tree:

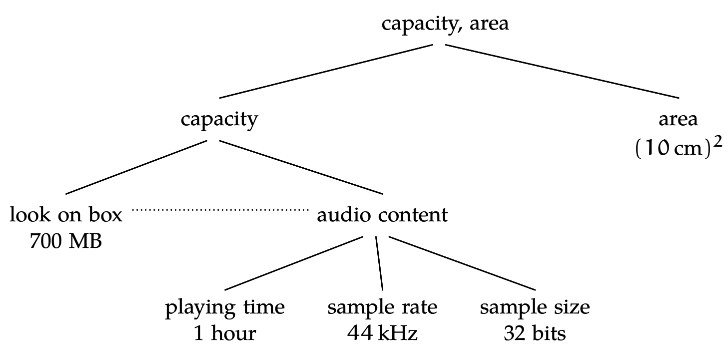

The entire paragraph flows well, moving from leafs to more branches. It seems as though some of the more important vocabulary is very abruptly introduced. I would prefer reading about the idea and then the name.

what is the difference between sample rate and sample size?

I think that the area leaf should have two branches: Estimate and Measure. After all, if you can look at the box to get the capacity, then you should be able to measure the size

these trees would definitely be helpful in explaining what you're doing, to someone else.

this is much easier to follow than the previous sections. i like diagrams that show where we're going.

I agree with the above - the estimations made in the previous statement make much more sense now, in terms of how they go together.

I also really like this tree, and even though you don't explain trees until this section I think it might be nice to see some representation of this earlier on.

Agreed - it might be nice to place these trees in the previous sections to make the reasoning/logic more clear.

could this section even go first? Introduce this as the first section in order to say, "this is how we will solve this problem", and then use the earlier sections to fill in the numbers?

I agree this is a much easier way to understand the concept of divide and conquer. I'm looking forward to the next section to see another example to see how it introduces an idea as opposed to reviewing it.

I agree with everyone, this tree is an excellent example of divide and conquer and illustrates the concept well. I think I would have had much fewer problems with the earlier sections had this diagram and tree explanation been put earlier in the text.

When does dividing become unhelpful and cumbersome? How many methods should we stick to and how detailed should we "divide" them?

so we are no longer at two different portions of the problem now, but rather two different methods solving the same issue?

CD. The second method subdivided into three estimates: for the playing time, sample rate, and sample size. Accordingly, the 'capacity' node sprouts new branches – and a new connector:

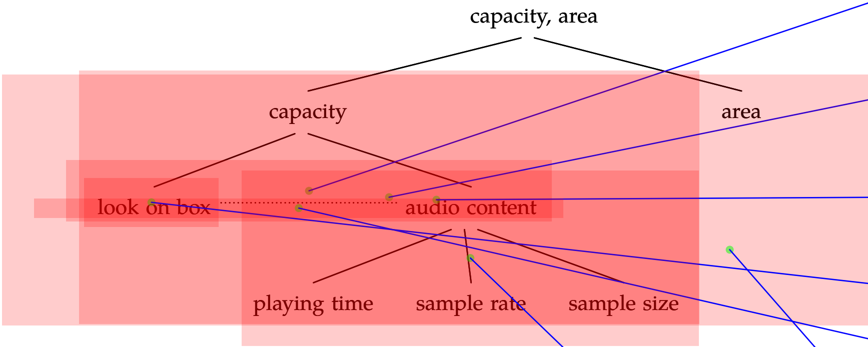

The dotted horizontal line indicates that its endpoints redundantly evaluate their common parent (see Section 1.2). Just as a crossbar strengthens a structure, the crossing line indicates the extra reliability of an estimate based on redundant methods.

The next step in representing the estimate is to include estimates at the five leaves:

1. capacity on a box of CDROM's: 700 MB;
2. playing time: roughly one hour;
3. sampling rate: 44 kHz;
4. sample size: 32 bits;
5. area: $(10\,\text{cm})^2$.

Here is the quantified tree:

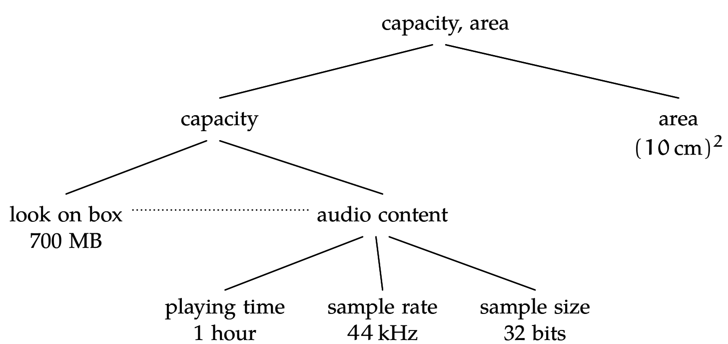

the dotted line is a little confusing. is there another way to represent intelligent redundancy?

Looking at this diagram now is helping me better understand this half of the approximation from the previous chapters. I feel as though this isn't really a separate method, but more of a visual for the previous methods. It's not like this tree diagram really replaces what we had to do before.

This seems like a confusing system – in the top level, the two branches are both necessary to estimate the pit spacing, yet in the second level, either looking at the box or calculating the audio content is necessary. Perhaps we need to distinguish between ANDs and ORs?

"Look on box" doesn't seem to fit the way the other branches are named. This is an action whereas the others are not. Perhaps use "labeled capacity" or "manufacturer's capacity."

Is the dashed line used in refer to redundancy?

Question answered in the next sentence. Oops. ;)

This comment box made the dashed line hard to see, and I didn't notice it until I read the comment. Does anyone know if there is a way to move the comment box outline?

I found you can toggle the visibility of the comment boxes by clicking the little icon of a folder with a paperclip at the top of the Nav tab on the left of the screen.

This graphical method provides a very good way to think about approaching approximation. But are there other aspects we can incorporate into the tree like the soft sides of the calculations?

I'm not convinced that this tree diagraming is consistent. Shouldn't "Playing time, sample rate, sample size" be one node off of audio content (as a method for estimating the audio content), and then playing time, sample rate, and sample size be leaves off of that node? It should be analogous to how "capacity, area" is a leaf off of pit spacing, later, and capacity and area are branches off of that. Because what if we had another method to estimate audio content?

I agree with this, I am a little confused with the notation. I think the bar and leaves should almost be reversed - The bar connecting things that interact to form the parent, perhaps with the operations in which they act detailed on the bar. Separate "leaves" can then describe redundancy. This way things that interact are always connected.

CD. The second method subdivided into three estimates: for the playing time, sample rate, and sample size. Accordingly, the 'capacity' node sprouts new branches – and a new connector:

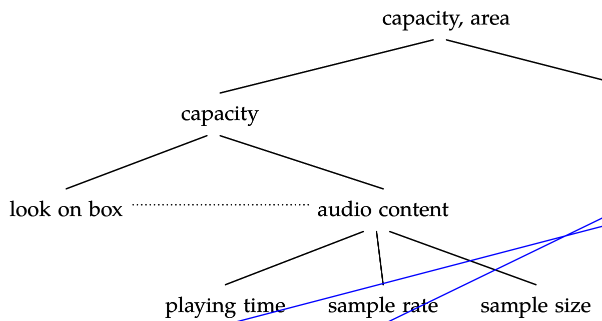

The dotted horizontal line indicates that its endpoints redundantly evaluate their common parent (see Section 1.2). Just as a crossbar strengthens a structure, the crossing line indicates the extra reliability of an estimate based on redundant methods.

The next step in representing the estimate is to include estimates at the five leaves:

1. capacity on a box of CDROM's: 700 MB;
2. playing time: roughly one hour;
3. sampling rate: 44 kHz;
4. sample size: 32 bits;
5. area: $(10\,\text{cm})^2$.

Here is the quantified tree:

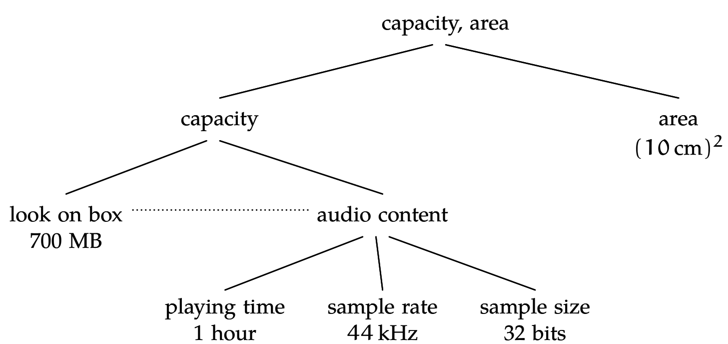

There should be some way to explain how these are put together to find the 'audio content'

does this mean that because we can look on the cd and read how much audio time is on the box, we can figure out how big it is, or because the size of the cd dictates how much audio time is on the cd?

The line is a bit hard to see now with everyone placing comments on the tree.

So a horizontal bar means two different methods of finding the same solution, while if there is no bar both children are necessary to come to the conclusion (the parent)?

Right. I have used tree notation to break down problems before, but I think this idea of representing redundancy on there too is flawed. A dotted line doesn't indicate OR to me.

A related question: Say we had a fourth node under "audio content" and it had a dotted line to "sample size", do we assume that the fourth node is a redundant method for calculating "audio content" on its own, or simply an alternative to "sample size"?

I _think_ it would be an alternative to sample size, whereas multiple dotted lines, or perhaps a different notation, could describe an alternative method of deriving audio content.

If you wanted to represent alternatives to calculating sample size, would you have two branches coming off of the sample size node indicating the different methods?

I think I am confused because it seems like the branches represent two different things – how to divide the main problem into subproblems, and also how to approach the problem in different ways. (I think this is the same confusion from the last section about how redudancy relates to divide and conquer)

I agree. It seems like the tree structure is trying to accomplish too much at once. However, it is also dependent on how deeply we are to use these examples (whether they are to be robust to every situation or are to just designed to convey a specific concept)

Why does it not represent "or"? Rather, if the point is that it is supposed to be redundant - it shouldn't be that you have to choose an "or" but keep in mind both answers to see the robustness of your conclusion. Also, I think it makes more sense to draw 2 branches coming off of sample size, as that is what we did for when we had 2 estimations for capacity.

Also, using dotted lines can get messy when making these tree diagrams by hand

CD. The second method subdivided into three estimates: for the playing time, sample rate, and sample size. Accordingly, the 'capacity' node sprouts new branches – and a new connector:

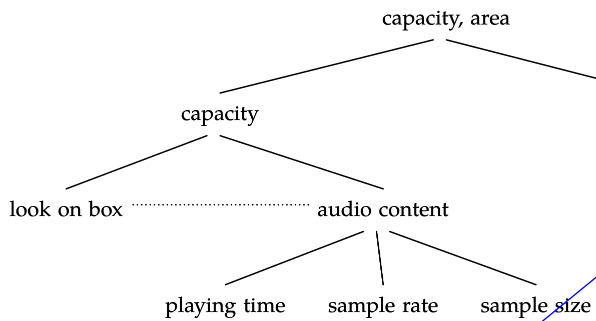

The dotted horizontal line indicates that its endpoints redundantly evaluate their common parent (see Section 1.2). Just as a crossbar strengthens a structure, the crossing line indicates the extra reliability of an estimate based on redundant methods.

The next step in representing the estimate is to include estimates at the five leaves:

1. capacity on a box of CDROM's: 700 MB;
2. playing time: roughly one hour;
3. sampling rate: 44 kHz;
4. sample size: 32 bits;
5. area: $(10\,\mathrm{cm})^2$.

Here is the quantified tree:

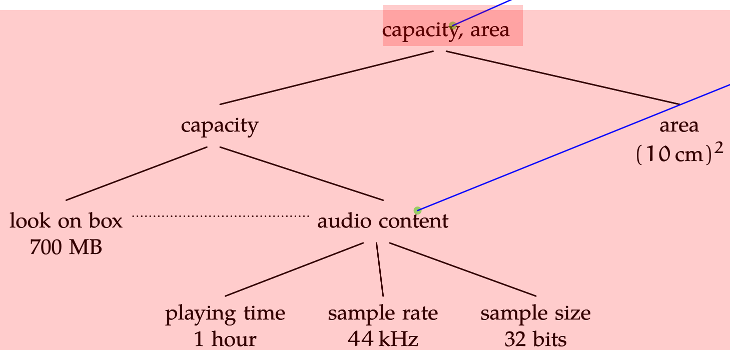

I think the dotted horizontal line is actually making the diagram more confusing.

But this isn't meant to imply that they depend on one another, correct?

Should we also include +/- tolerance values here in the tree so that glancing at the tree gives us an idea of our error?

I think so. If nothing else, this will give us a way to assess the value of the different estimates that we came up with through the different redundant methods

I like this list. It might be helpful to start with a list of things like this even if you don't have a method planned out yet.

I also think this list is nice; at least in the text. it gives a good layout of the 'tree'.

I think everyone wishes they had this list at the beginning, but then it wouldn't be an estimation. I think I'm beginning to understand that a big part of estimation is just determining which method requires which information for an estimation.

I agree with this above statement, there are so many relevant statistics but determining what is useful is a very crucial part and the tree allows us to start doing this without putting an excessive amount of thought into it

Since audio content depended on some function with playing time, sample freq, and sample time as variables, shouldn't the top of the tree be pit spacing, and then that would branch into capacity and area, since pit spacing is a function of capacity and area?

is there any other way to word "capacity, area". this is around the third or fourth time i've seen this and it still looks as it did awkward as the first time i saw it.

I always have trouble looking through my notes later. I'm going to start using trees when I do the hw problems

CD. The second method subdivided into three estimates: for the playing time, sample rate, and sample size. Accordingly, the 'capacity' node sprouts new branches – and a new connector:

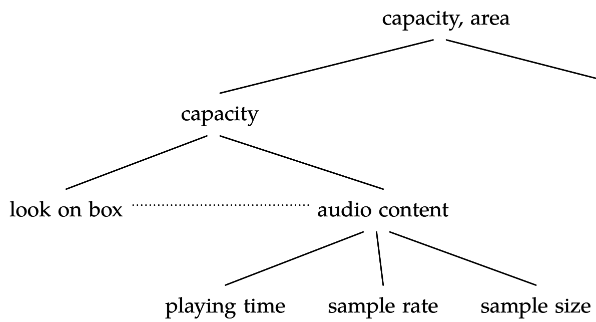

The dotted horizontal line indicates that its endpoints redundantly evaluate their common parent (see Section 1.2). Just as a crossbar strengthens a structure, the crossing line indicates the extra reliability of an estimate based on redundant methods.

The next step in representing the estimate is to include estimates at the five leaves:

1. capacity on a box of CDROM's: 700 MB;
2. playing time: roughly one hour;
3. sampling rate: 44 kHz;
4. sample size: 32 bits;
5. area: $(10\,\text{cm})^2$.

Here is the quantified tree:

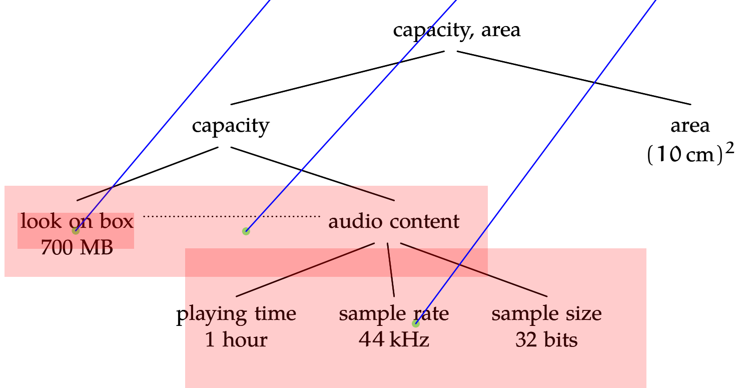

I realize the redundancy is to strengthen the estimate, however, given that one of the values (the 700MB) is a "given" (found on the box) is the other one really necessary? I understand its point if both values are estimates (again, redundancy) but if one is a given value, why would this matter?

I completely agree with that

instead of thinking of it redundancy to check, could we instead just treat it as having multiple methods to tackle the problem. for example, if the box is lost.

Is this being over-redundant, or in general should we not trust this kind of data from the supplier?

This tree is really helpful but I still don't understand how I would know that audio content breaks down into these 3 components. In a problem, would this be defined for us or is this something that would be regarded as common knowledge?

Yeah, this I don't think I'd split that into 3 different leaves.

It may not be common knowledge before reading the section, but hopefully afterwards it is!

To answer more specifically: It depends on the level of the question. An easier question, might give you the breakdown and ask you to estimate three numbers. A harder question, more like real life, might leave you to figure out any way that you can. Both questions have their value (one cannot play Bach without first learning to play scales).

The overall concern is a big reason for my grading system where correctness is not relevant. Don't worry about doing 'badly' on a question. Just keep your eyes on the prize: If you are diligent in this course and continue to practice, you'll get fluent doing all of this. So, just aim for that.

The final step is to propagate estimates upward, from children to parent, until reaching the root.

Draw the resulting tree.

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly 1.4 µm.

Propagating estimates from leaf to root gives the following tree:

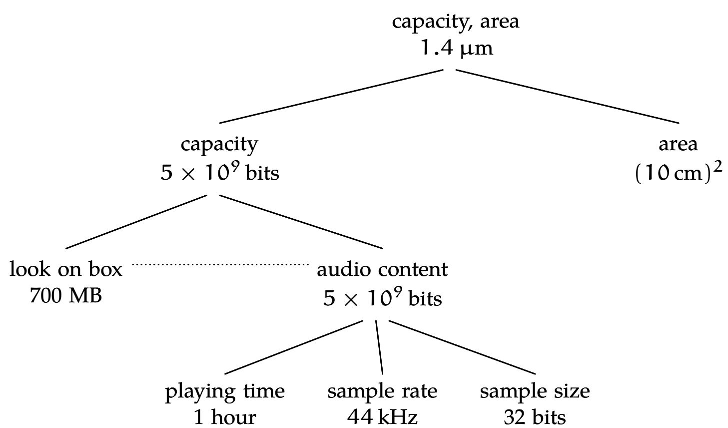

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

Draw a tree that includes these methods.

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends

Why change the vocab to children and parent? Is this a more general way of explaining trees or just an analogy?

The vocabulary of children and parent is widely-used and accepted when referring to a tree data representation. Although I do agree, the change from root/node/leaf to parent/child is odd.

I agree it is confusing, but it's okay here. In fact, it's a good way to introduce vocabulary.

This note is very well organized and easy to follow.

What do you mean by this? You already gave us the tree...or do you mean, 'fill in the resulting 'non-leaf' nodes?

I would like to see this somewhere on the tree, even if it was just an 'x' between each of the 3 leaves to show that they are multiplied to get back to the earlier branch.

it gets too cluttered that way, and you lose the benefits of having a tree in the first place.

maybe instead of drawing the symbols between the leaves, making the tree more cluttered, it would be useful to draw the operation symbols between the branches

I agree–some way of seeing how the branches combine to form the parent node at each level has been helpful to me when I draw trees. I think the tree is most helpful for conceiving of the dependency between parameters, and nearby or separate equations can explain the mathematical relations between them.

The relationship might not always be multiplication. And we might not know what the relationship is. But using the tree, now that we know what the dependent parameters are, its easier to formulate a relationship.

The final step is to propagate estimates upward, from children to parent, until reaching the root.

▶ *Draw the resulting tree.*

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly 1.4 µm.

Propagating estimates from leaf to root gives the following tree:

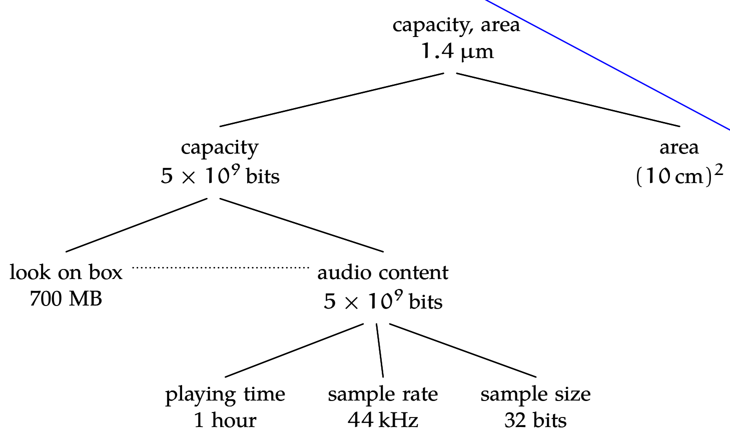

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

▶ *Draw a tree that includes these methods.*

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends

how do you know it's the product of these three variables? Is there a formula for this?

I think it's product as in derived from separately, not actually all three numbers multiplied together.

I think it's that the way they set up this tree, it happens that playing time*sample rate*sample size gives audio content in bits if you do unit conversions for Hz and hours–in the first reading this was how they did it too, it's just much more clearer now using the tree

Exactly, multiplying hr*KHz*bits=bits, thus we used unit analysis (dimensional analysis).

ok- but this is not the general case, right? like multiplying sibling leaves isn't always the operation necessary. Is there a way to represent the operation performed on a set of leaves in the tree?

I don't think that is the general case. I think depending on the problem, the necessary relations will vary. The tree does a good job of breaking down the problem into manageable pieces (divide and conquer) however, it does not provide much insight into how to derive a solution from all of these broken up pieces. Unless someone else has seen this. thoughts?

perhaps there could be a legend on the side, it could clearly state that the dotted horizontal lines show redundancy

It seems strange that they found exactly the same answer. The numbers 44hz and 32 bit sampling seem too accurate to have been guesses, unless the guesser happens to know that they are the real values, I don't think the average student would be able to find this answer with such accuracy.

The final step is to propagate estimates upward, from children to parent, until reaching the root.

▶ *Draw the resulting tree.*

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly 1.4 µm.

Propagating estimates from leaf to root gives the following tree:

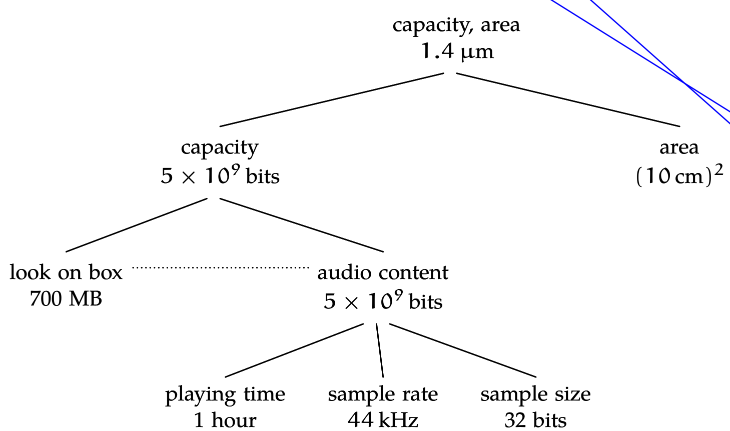

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

▶ *Draw a tree that includes these methods.*

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends

If they didn't agree, how would you represent this on the tree?

by erasing the wrong one? (kidding) If they didn't agree closely enough (enough being a key word), then you would probably have to reevaluate to see which one is in error, or find a third method to serve as a tiebreaker?

Hopefully they wouldn't vary by too much. However, we may often have an implicit reliability with each method. Surely the box is more reliable than our estimation via pits. In which case, we would have to use our gut to determine if we choose one over the other or pick a value in the middle closer to one than the other.

Would this be where we could add in the tolerances to our estimate? If the look on box and audio content vary from each other, could we pick a value between the two and set a tolerance based on how far apart they are?

Some people may not understand how 700MB = 5*10^9 bits. Maybe a quick demonstration or mention of the conversion factor may help to make this material clear to the non-course 6 students?

I agree with this I had to refresh myself with this conversion as well, and it might be nice just to see it there instead of looking it up.

I just learned it so I will post it here. Bascially, to convert megabytes to megabits, you multiple by 8. Then to convert megabits to bits, you multiply it by 1048576. Thus 8*1048576 results in 5.8E9, or approx 5E9.

1,048,576 = 2^20

I think a bit of background of bits and maybe even a section on estimating binary representations of numbers could be quite interesting for this class

these formulas should probably be included on the tree somewhere. It would make the entire tree much easier to read and follow.

It would be nice if to the left or right of the tree, the different equations used to go up a node were shown

I agree, the tree shows what are the necessary calculations/number needed but does not indicate how they are necessarily used. In the example here it's not a big deal but I'd imagine in more complex examples the formula written into the tree would be handy

The final step is to propagate estimates upward, from children to parent, until reaching the root.

▶ *Draw the resulting tree.*

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly 1.4 µm.

Propagating estimates from leaf to root gives the following tree:

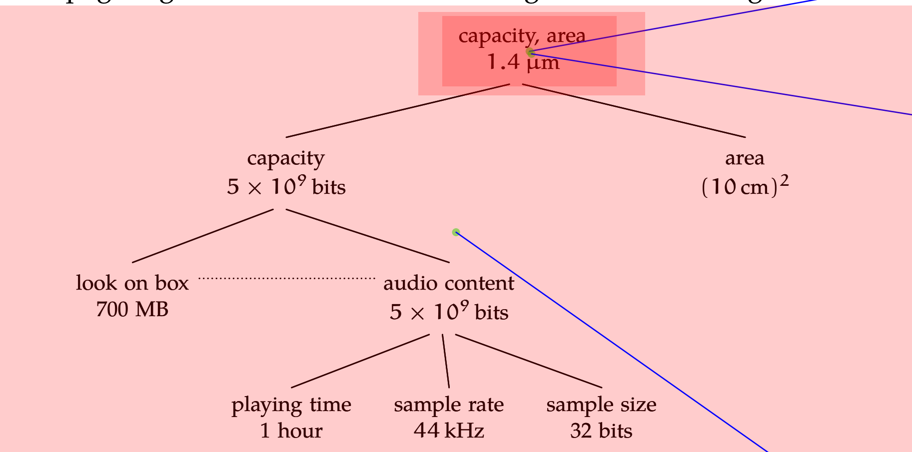

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

▶ *Draw a tree that includes these methods.*

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends

I believe this is a large problem that may need to be looked at (or hopefully we will look at later), and it is that this does not seem like common knowledge. Although we can find all of this information, where can we get these equations from? Can we better estimate it if we don't know? I understand the concept that we may be doing estimations more in our field, but what if someone asks me this, and I'm not CS, and they want a legitimate answer? I feel like this divide an conquer method is effective in broad cases, like maybe the example on the next memo, but in technical cases, I can't see a strong benefit, though I like the idea of it.

This is a very easy to follow and clear demo of divide and conquer. Is there a way to incorporate/tell how to mathematically process the tree (i.e. how to know whether to multiply, divide, add etc)?

It also might be helpful to have the variable we are solving for (pit spacing) at the top of the tree above capacity, area. I'd like to see it there just as a reminder of what the end product is.

Oh didn't see this comment - I posted a new comment with the exact same opinion. I think this would make the formatting of the tree more consistent.

Given that the number we place here is the spacing in between pits, shouldn't the tree's top level be "pit spacing" instead "capacity, area"?

I believe the point is that the number given is the pit spacing as calculated using the capacity, area method.

However, I agree with you and all the others that perhaps it would be best to start the initial introductory tree with "pit spacing" because it has caused so much confusion so far, and it is clear that most people understand the final tree and are just getting hung up on the fact that we are introduced to a tree that doesn't start at the true root of the grand final tree.

The whole tree is shown on the next page, I think that clears it up.

oh, yeah, the next tree clears it up a bit more, perhaps having a legend would clear up more confusion since some same height leaves have a different relationship than other same height leaves.

This is definitely helpful if you don't get all the way through the problem and need to come back to it latter.

The final step is to propagate estimates upward, from children to parent, until reaching the root.

Draw the resulting tree.

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly 1.4 µm.

Propagating estimates from leaf to root gives the following tree:

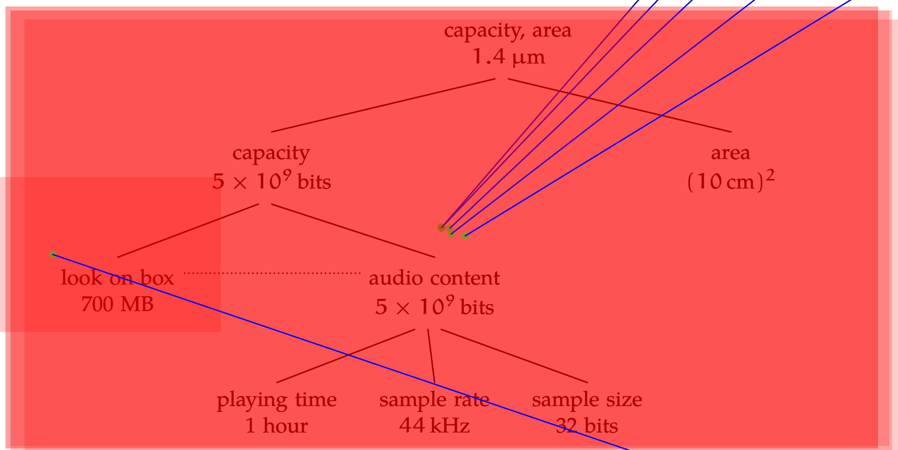

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

Draw a tree that includes these methods.

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends

These charts are very helpful, maybe they should come in the earlier sections?

These charts are very helpful, maybe they should come in the earlier sections?

I agree that it's difficult without the various operators so it would be nice to see on the lines what the mathematical relationship is.

I thought the example of MIT's budget today in class left a little more to the imagination on the tree's and was more open....u can get stuck into estimating many branches and going deep down in the tree...when in reality trees should be a fast way to do the analysis

Looking at this tree the original person can understand it, but without the formulas no one else could

Definitely true. Although this makes the concept of how the estimations are broken down clearer, it gives no indication of how to put the same estimations back together and find the answer.

Maybe the "branches" should include a process as to how the answer was reached from the given information. This would make the lines more literal pathways between raw information and the conclusions drawn from it.

I agree that it would be useful, although I think the point of the tree is just to concisely breakdown the steps and numbers and leave out detailed explanations

I also agree. Simple equations could be added between branches to indicate how the solution above was generated. Looking at the chart, it isn't obvious that playing time*sample rate*sample size=audio content.

I was confused at first, but after you study the tree a little it does become quite apparent. Realize that these numbers aren't going to be trival to many people and they are going to have to think about what is going on.

The tree is made to show all the pieces needed to solve the complex problem. It is only a way to sort the info. I think it was a stretch for the author to say that you propogate estimates from parent to child so matter of factly.

So where does it show whether you chose to use the audio content data, or the look on box data?

more compact, yes, though inferior in explanation, right?

The final step is to propagate estimates upward, from children to parent, until reaching the root.

▶ *Draw the resulting tree.*

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly 1.4 µm.

Propagating estimates from leaf to root gives the following tree:

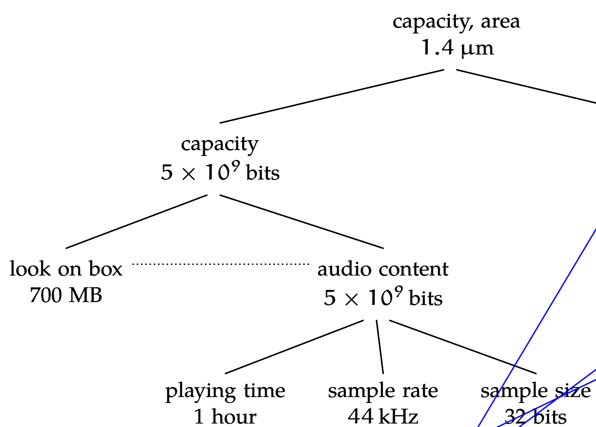

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

▶ *Draw a tree that includes these methods.*

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends

As a final way to put together all of the information, it would be nice to see the all of the calculations multiplied together to form the result.

Definitely. The tree itself is a great summary, given that we're reading this over the span of a few days, and perhaps the actual calculations (with units) would help reinforce what we've been doing.

i'm a little confused here. is the tree supposed to be a "new" way to solve these problems? because it seems like just a different visual representation of the same approach.

I think the tree is simply a new way to display the information and how the approximation is split up into different components.

I couldn't agree more, which is why I think this should be introduced with the material from the beginning to make it easier to understand.

I agree with this sentence and I do believe trees are a lot clearer. They also give the person who examines trees time to figure out how you got the answer by themselves. Which is definitely an improvement on learning rather than just believing a answer given by a professor must be correct.

How much variation is acceptable between multiple trees for using the same method of estimation?

It seems like we could help account for variation (horizontally), and make the robustness clearer, by adding the (+/-b) to each of the leaves. (I am referring to the estimated accuracy of each value's magnitude as displayed in the exponent.) Though, this may just have the effect of cluttering the tree diagram...

By the end of this section I've forgotten what we are trying to approximate in the first place. It would nice to have a reminder after we have done all the little estimations.

agreed.

The final step is to propagate estimates upward, from children to parent, until reaching the root.

▶ *Draw the resulting tree.*

Here are estimates for the nonleaf nodes:

1. *audio content.* It is the product of playing time, sample rate, and sample size: $5 \cdot 10^9$ bits.
2. *capacity.* The look-on-box and audio-content methods agree on the capacity: $5 \cdot 10^9$ bits.
3. *pit spacing computed from capacity and area.* At last, the root node! The pit spacing is $\sqrt{A/N}$, where A is the area and N is the capacity. The spacing, using that formula, is roughly 1.4 µm.

Propagating estimates from leaf to root gives the following tree:

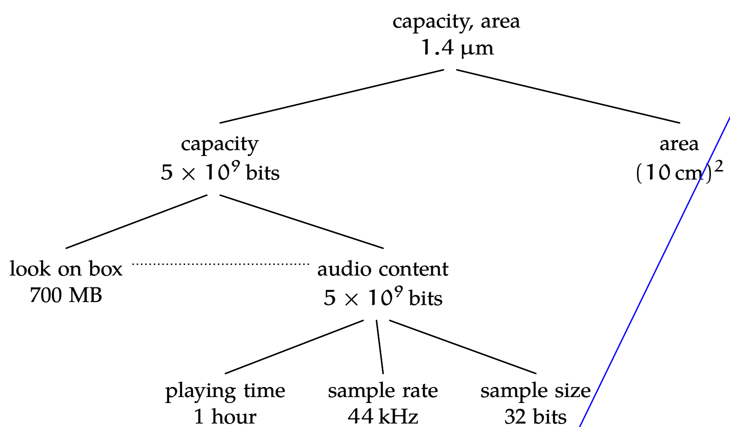

This tree is far more compact than the sentences, equations, and paragraphs of the original analysis in Section 1.1. The comparison becomes even stronger by including the alternative estimation methods in Section 1.2: (1) the wavelength of the internal laser, and (2) diffraction to explain the shimmering colors of a CD.

▶ *Draw a tree that includes these methods.*

The wavelength method depends on just quantity, the wavelength of the laser, so its tree has just that one node. The diffraction method depends

This sentence reads awkwardly- is it supposed to say "depends on just ONE quantity"?

good catch, probably!

I also got tripped up here. I might remove the word quantity altogether and just state "depends on the wavelength of the light"

(The sentence needs "ONE", as noted in this thread.) I remember debating this suggestion mentally when I first wrote the sentence. The reason I didn't phrase it that way is that I want to emphasize the pattern that one quantity means one node. If I mention only the wavelength, it's harder to generalize from this example to the pattern.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

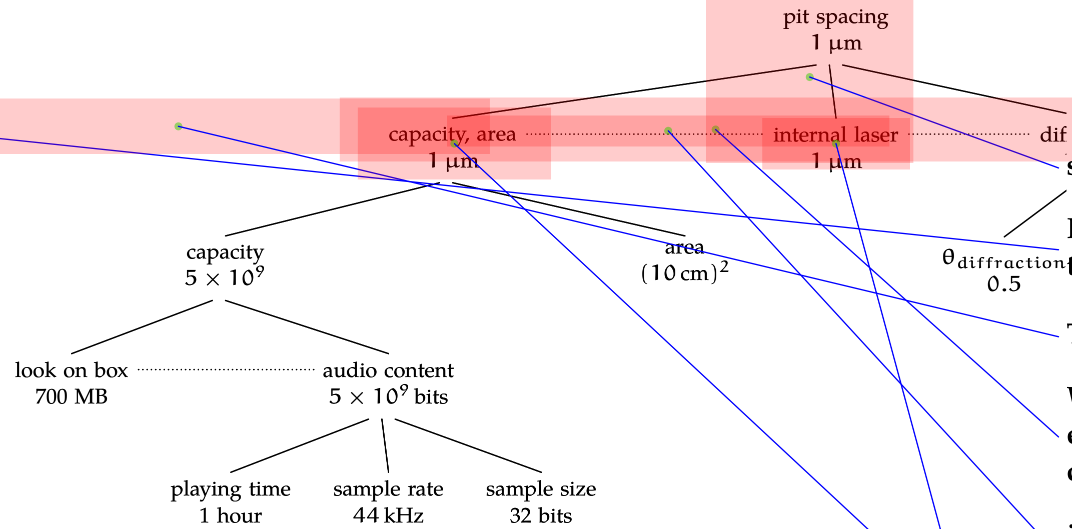

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

In my experience, sometimes I may need to come up with individual trees as I get information, then once I have everything, i put all the pieces together like a puzzle. I'm not sure if thats the best way to do it, but thats the best method I know of.

Yeah, I agree–because we don't always have our whole plan written out. I guess in this case it would work, since as we move up the branches, everything is simply multiplied together, so separate trees made earlier could be multiplied together

should this larger tree be written down before we begin to crunch numbers at all?

I think that it's ok to use 'capacity,area' here because there are other methods, however the other tree should still have pit spacing because that is what we're trying to find.

This should be 1.4 because that it what you found it to be on the previous page

Would it make sense to arrange these in order of easiest to hardest to work out, or rough estimate to more known quantities? For example, the internal laser is a rough guess, while diffraction and capacity/area are much cleaner

it would seem the redundancy line refers to data that has the same value regardless of dependency. Clarified from earlier

So I made a comment about this section earlier saying that this should say "pit spacing" but now I see why it doesn't - I'm no longer quite sure how to address this but something should be done to make it slightly less confusing. At first glance, relative to the rest of the tree, I would think that this section means "capacity, area" = 1um, which makes no sense.

I was unclear about this also. I thought it should have been pit spacing like you. Seeing as it is not, though, I'm confused as to why the "capacity,area" node needs to be there. Couldn't "capacity" and "area" simply be direct children of "pit spacing"?

I like this one, nice and simple.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

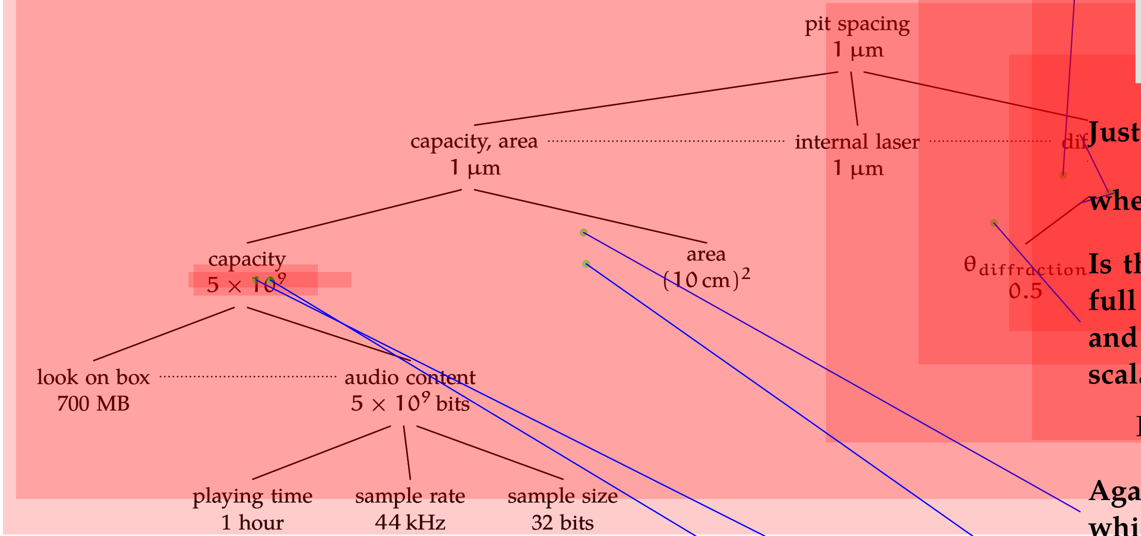

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

i'm assuming this just got cut off by the margins?

You're right. (It used to fit on the page when it was typeset on 8.5x11, but even then it was beyond the margin). I haven't yet figured out how to redraw the tree to fit in the margin and still be clear.

Just a minor comment: this area of the tree got cut off when I printed it.

where did this part of the tree go?

Is this a formatting error or is there some other reason that the tree cuts off without the full diffraction example shown? I'd like to be able to see the two children of diffraction and how they actually interact together - if its the multiplication thing, a division with a scalar or sqrt...?

I believe Sanjoy answers that a couple of questions above you. Formatting error.

Again, we really seem to need a way to distinguish redundancies from multiple branches which are all required.

Can anyone see this entire tree? I can imagine what should be on the other side but I'm just wondering if this is cut out for everyone.

units = bits

Why don't you keep the units on this version of the tree? Shouldn't you indicate bits?

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

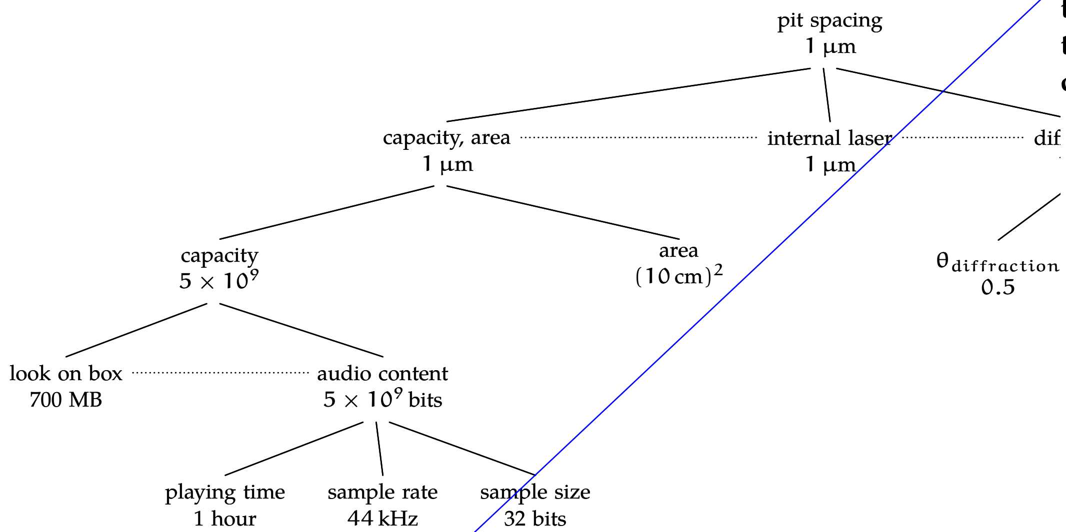

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

The tree leaves me wanting to know what the equations are that lead up from the leaves back to the roots. There must be a way to include an equation along the sides of the branches to show how on goes from sample slze back to audio content, especially for these calculations that contained a lot of new information for most of the class. While this might make for a nice graphic summary, I feel like it doesn't really have that much content in it.

That's an interesting idea. The most common equation for each node is "multiply all the children to get the parent." But not always. For example, pit spacing is not capacity * area but sqrt(area/capacity).

Even so, the operation is usually multiplication of the children after raising each child to a particular exponent. Acompact representation for that would be to place the exponent along the line connecting the parent and the child. For example, the line to the capacity node (that says 5*10^9 bits) would have -1/2 on it. Similarly, the line ot the area node would have 1/2 on it.

In cases where that framework is not general enough (e.g. reaction rate might be e^-1/T), one could give the explicit formula at the node.

What do you think?

That sounds like it would work pretty well, as long as there was a table of symbols or something that reminded the reader what the 1/2s meant.

That sounds like a simple way to include information about how errors in the estimations for values lower-down in the tree would propagate upwards. It would also lack any constant coefficients present in most equations that are a product of variables, but that's not really relevant to error propagation.

I like the idea of having the equations at the nodes; I'd say that it probably defeats some of the purpose of the estimation, but if you have no equation (no previous knowledge), then whether you can estimate some numbers doesn't really matter. I can agree with that,

Yes, I think that would be a useful, simple, addition to the tree structure.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

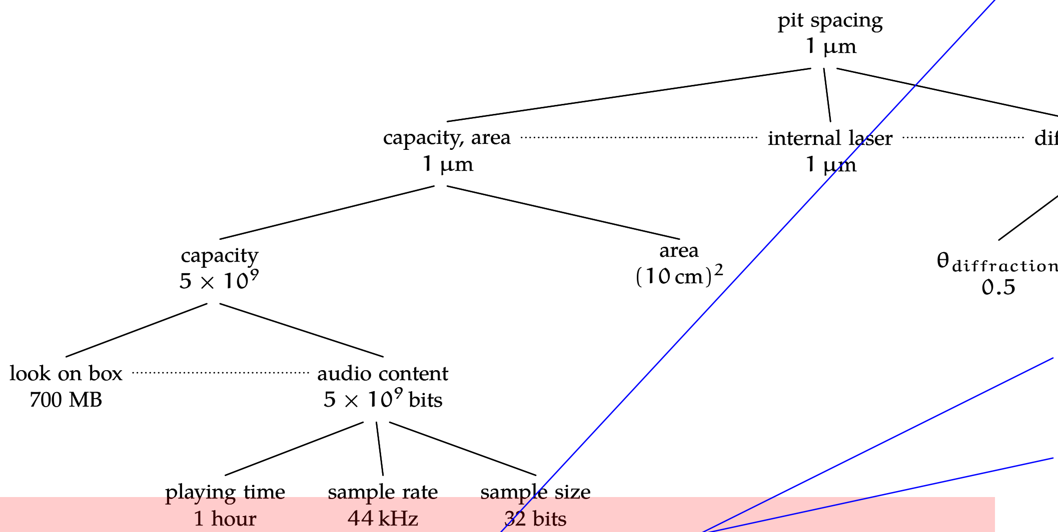

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

I agree that this visualization makes everything much easier to understand, so perhaps it could have been used earlier on and integrated into Sections 1.1 and 1.2 instead of having a separate part by itself?

I agree that the visualization in the form of a tree is very simple, making the divide and conquer method more clear. However I disagree that it should be shown in the previous two sections because it is a good way to summarize the previous approximations and bring everything together.

I also think that the tree is placed best as a summary. While it would be helpful to have this tree from the start, we should be allowed to go through the mental process of creating our own version of this tree ourselves so that when it is shown later (like in section 1.2) we can match it with what we had envisioned

typo: makes

Overall, I think some of the less obvious points, such as the operations being done, could be put onto the tree. Having said that, the tree method works very to show exactly how the answer was found in a very systematic way.

Exactly. This is why I think it should come first, it is much less complex and easy to understand. Then you can elaborate into each branch in the next sections.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

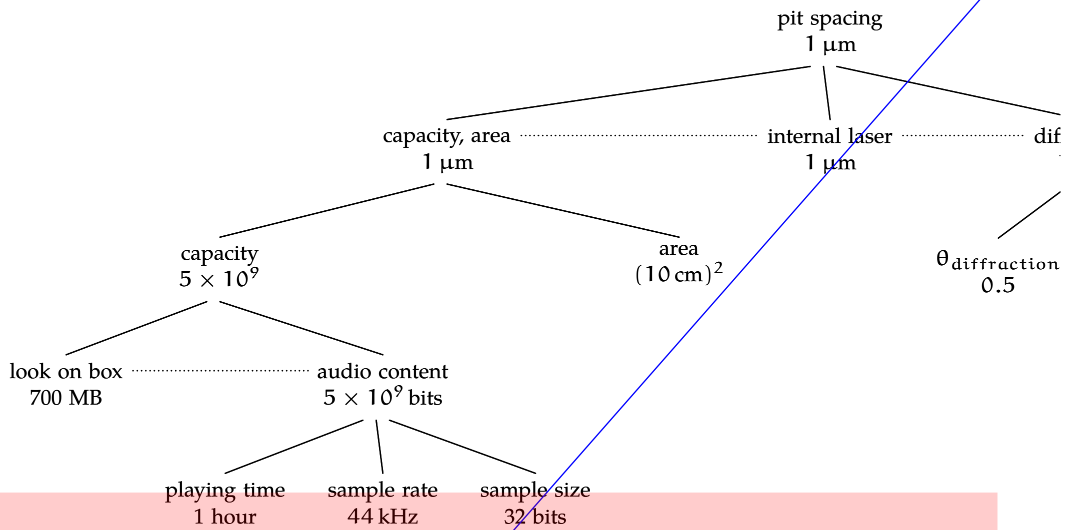

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

How much oil does the United States import (in barrels per year)?

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

For large trees, is there a strategy of what to expand first? (i.e most complex branch, branch in which you think there are most layers_)

I have a similar question. from what it seems, the tree diagram is basically an easier way to visualize the estimations we are making to arrive at our one large estimation. But the tree doesn't do much more–we can't get from it any new insight as to where to start or how to calculate certain things. So even with trees, we're still pretty much on our own in terms of figuring out what values to estimate and multiply together, etc.

It seems to me that it would be most beneficial to expand the branch with the most numerical calculations, and double check your answer with more other practical methods such as diffraction.

I was wondering about this also. When doing these problems, I usually don't have it so mapped out in my head. The trees are super useful in recalling the process, but I'm not sure how much they would help me solve it.

It seems to me that the tree should guide your estimations. If you reach a leaf which seems difficult to estimate immediately, then look for a way to divide it into a set of easier calculations. Once you've reduced your problem to a set of simpler calculations then work backwards to estimate your desired quantity.

The tree is a great visualization tool, however I like knowing where all the numbers came from. The analysis in Section 1 should be introduced after the tree because it would sink in faster with the picture of the tree in your mind.

I agree entirely. It makes the series of steps taken to solve the problem much easier to understand and follow.

If we had made trees throughout the sections, thus building up to this intricate and multi-tiered tree, I would have been able to synthesize the previous sections as we completed them. Making these connections makes it very clear why this is a method of divide and conquer and why their is intelligent redundancy.

I thought this entire section was very well written. Simple and concise. I can't think of any ways to improve it.

This was easy to follow, but at the same time, I don't know if I could just come up with all of this on my own.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

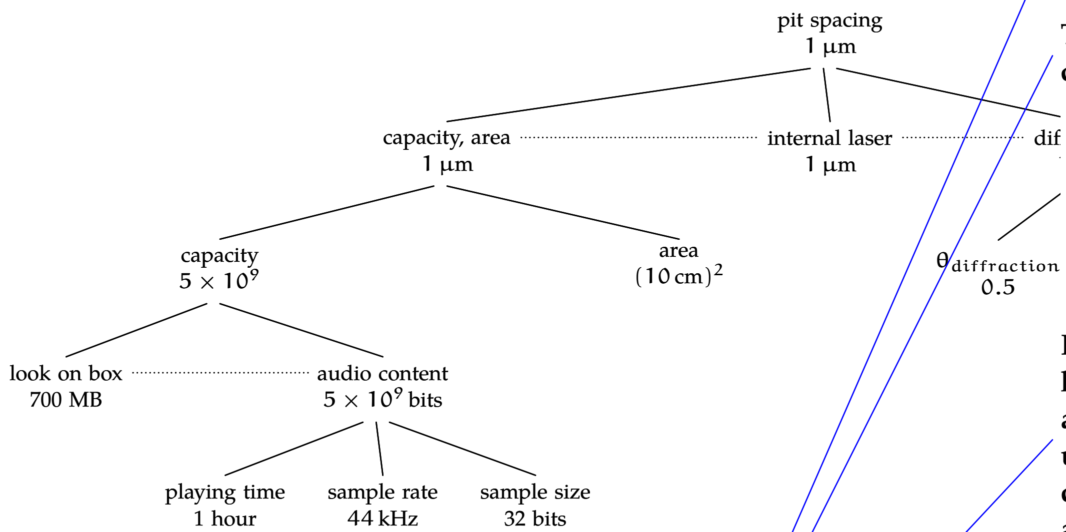

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

The overview definitely helps and I feel shows the consensus among different methods and a similar conclusion - are there any cases where you see a slight difference and end up having to normalize it in a way?

This paragraph summarizes the only use for trees. Once you understand the problem, you can write it as a tree so you can go back and easily understand what you did before.

That negates the notion of drawing the tree as you go along in the problem. I agree with that - I would have been able to initially figure out the approach and organize it into a diagram. Trees are just a way to easily convey information once you know what you're doing.

*wouldn't oops

If you hadn't given me this method to subdivide the problem, I would have no idea on how to solve the problem! With the some help from Google. I found: 20 barrels of oil are used by an average car per yr, there 200million cars in the US, about 66% of US oil is used in transportation, and the US produces 2000million barrels/yr. My calculations come down to being: ((20 * 200M)/(.66))-2000M approximately equal to 4000M = 4billion barrels a year

When we use each method separately, should we automatically add a factor to compensate for the other uses of oil that we are not accounting for in our tree diagram?

I think that each of these quantities have its own branch in the tree diagram.

I thought this memo was straightforward, so there weren't many comments to make.

I thought that I put these comments on the page before, whoops. Over spring break i went through and tried to find any other readings that my comments didn't make it from my paper to the nb system. thanks.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

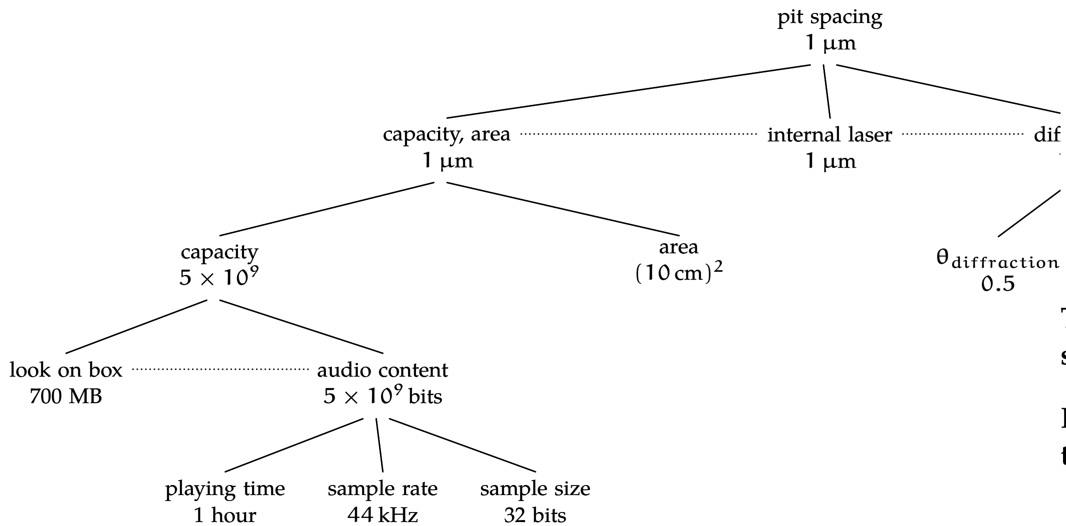

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

It seems a lot of estimation is pure luck. It happens that lowering 15,000 miles per car to 10,000 produced an accurate prediction. Isn't that luck more than skill? And if it skill what are tips for guessing accurately?

I feel like if you do that with all your numbers though, rounding up and down when working with your estimates. They end up canceling each other out, which is why I don't think luck has anything to do with it.

Using 15 would get you 4.5x10^9 barrels, which is still in that general range, especially if you use the method of calling it a "few." Maybe it takes luck or a good gut to get as accurate as he does but you can still get in the general range and at least be on the order of the correct answer, which seems to be what you're going for in these.

This seems like an overestimate, an energy report issued sometime last year or late 2008 said that we imported 60% of our oil, I imagine that rises yearly

I agree with previous comments, perhaps a better way to estimate N would be to look at the number of families in the US and assign 1 car/family

I see two problems with this approach 1) How do you define families/ Large chunks of the population is no longer part of traditional 2 parent families 2) Families tend to have more than one car

I supposed estimating the number of barrels of oil used by cars in the US is the easiest approach, but how would you go about the other nodes (other uses, fraction imported)? Clearly these are also important for intelligent redundancy...

I find it interesting and somewhat strange that in both class and this example we only did estimation for one of the quantities (cars here) and yet we reached a result that was equivalent to the total oil use/yr. Is the trick to largely overestimate in one quantity? ie. 3*10^8 cars in the United States. With these problems, should you always focus one quantity and overestimate it.

I feel that this example was a good way to demonstrate the concept, however, I feel that oil is more risky than what is presented here. Doesn't the price of oil fluctuate a lot? Or does the US just hedge for the price?

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

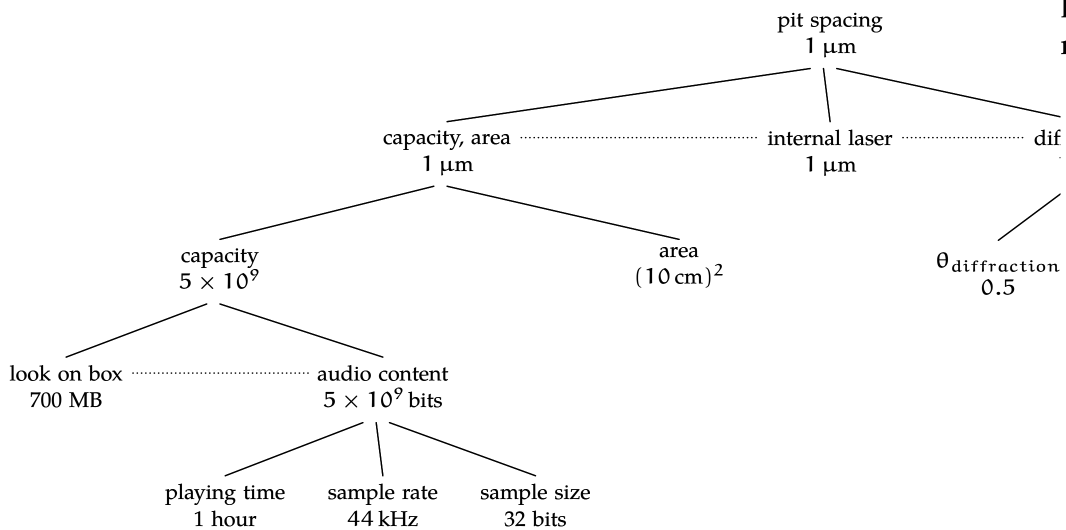

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

I think as a general comment, you should consider making the first examples for demonstrating a concept something that the reader probably already has intuition. It's much easier to learn a method if you have an idea for what the answer should be. If the problem is intuitive and easy, the reader only has to learn the method and not facts and other not as important things.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

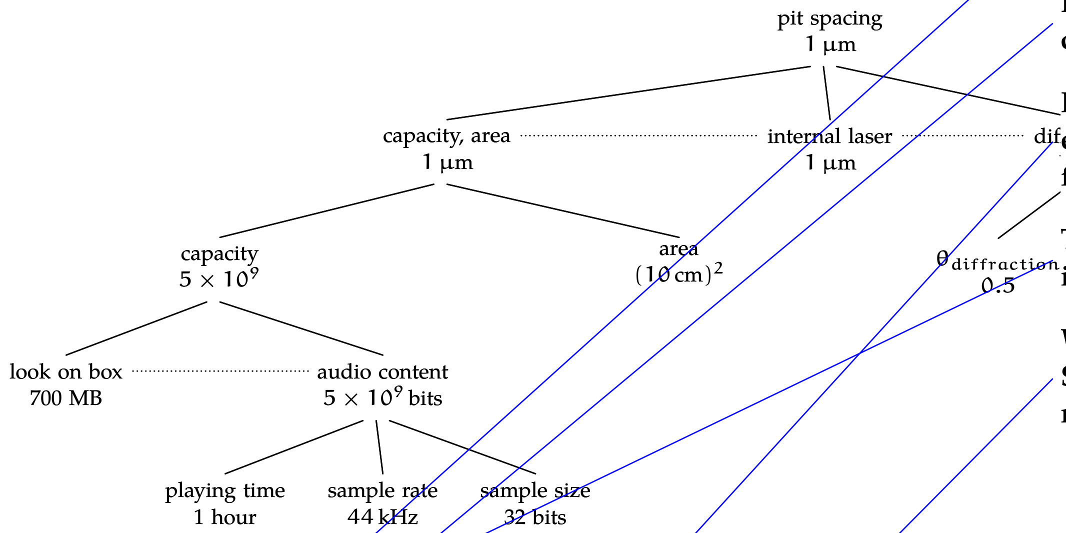

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

Read Section 1.4 for Friday (memo due Thursday at 10pm).

I'm actually having trouble coming up with a comment for this section. It seems quite clear and straightforward, and I think it helps clarify the technique well.

If you end up rewriting the first few sections of the book, I think this might be a good example / way of writing the whole divide and conquer section rather than just being a follow up to separate divide and conquer and tree sections.

These are similar to back of the envelope consulting interview questions. I like that this is helpful for not only in my engineering classes but job interviews as well.

Would a suitable alternative estimate be to calculate this as a proportion of the United States GDP (either by considering input costs for sustaining the economy's production, or relative to exports)? Or perhaps as a fraction of imports?

This just sounds like it would be much more confusing, and would require data that is even more obscure from the average person than the way that the article solves it, but it still might be worth trying if you know some of the necessary numbers, such as the US's GDP, or our total imports/exports. If you try it, do you get a similar answer to what the book does?

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

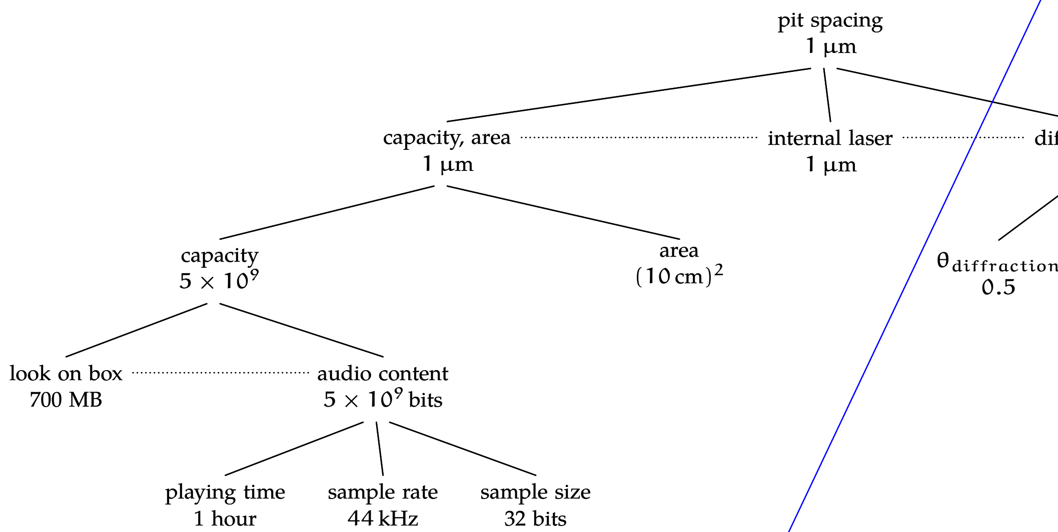

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

What if I don't have a sense of the volume of the unit (barrels)? I realize the conversions can be looked up, but why specify, since it's maybe more important that the answer be easy to grasp. The question could still ask for units of volume per year, but leave the volume up to the solver. This also makes the question more universal (we've already been switching around between American and metric systems).

I guess the thought was there needs to be some unit of measurement, otherwise you can end up with people who have answers in gallons, liters, crates...anything random. So if this is specified it makes everyone's answers closer, even though it leads to another level of estimation

Well that's all good and well, but it would have been better to pick one that people have a feel for, say liters.

I don't even know if barrel is necessarily a standard size. Maybe I interpret this barrel as being the barrel from the "barrel of monkeys" game.

Er, well, there's also been a lot about the cost of a barrel in the news, so it is a fairly standard unit for oil... I think it's fair to start assuming the reader knows at least something or can look it up.

Yeah, whenever you are talking about large amounts of oil, the unit is barrels, which equals 42 gallons.

I agree with the later answers - oil prices are always quoted in barrels. However, there could be a short description to give us an approximation of what a barrel is - even something like "approximately 40 gallons" in parenthesis would be helpful.

I agree. I think that what makes a divide and conquer problem easy is if we could somehow end up with branches at the bottom of the tree that relate to numbers that are accessible in our everyday lives, although it may not be the approach to take in all cases.

I think if you had to, you could estimate the volume of a barrel. Given that it's called a barrel, it would surprise me if its volume were less than that of a cube with side length of .5 meters, but it would also surprise me if it were more than that of a cube with side length of 1 meter. That happens to put the volume somewhere between 30 to 250 gal. The true value (42) ends up being in the lower end of that range.

Oops, he sort of goes into this later, using the cost of a barrel (if you know that...) and the price of gas.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

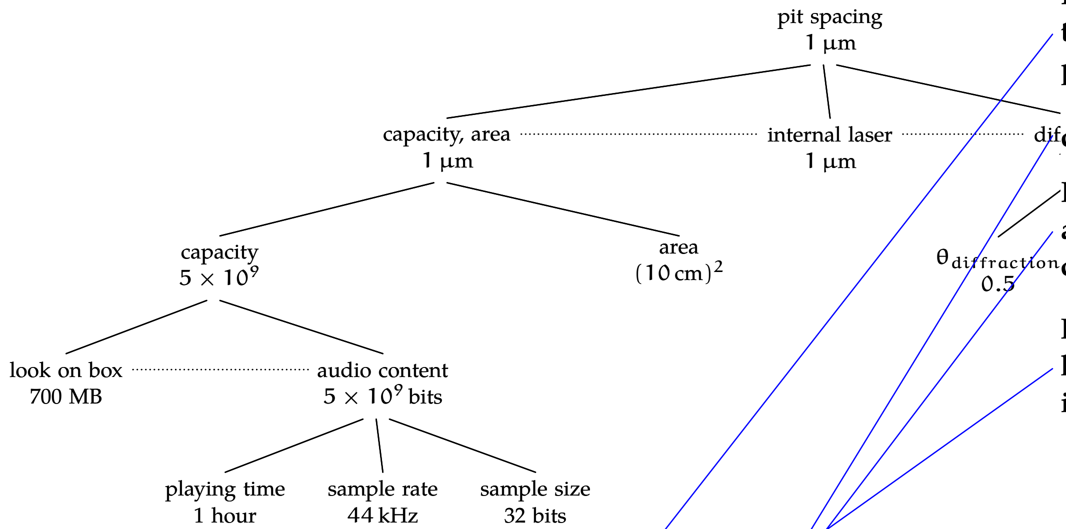

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

Oops, he sort of goes into this later, using the cost of a barrel (if you know that...) and the price of gas.

I just wanted to say I really like that you have this question in here. This number crunch type question comes up in consulting interviews a lot, and I'm really happy that it is a practical example I can relate to. Thanks for putting it in!

do we know how much oil is in a barrel?

It's seems to me that the first two quantities might be solvable but without any knowledge about the US oil industry it would be impossible to get rid of this value. I understand it can't be done without this value, at this point do we have to look up certain information?

For those of us who don't know the size of a barrel, perhaps here it could be added that how ever many gallons are in a barrel? I know it'd be easier for me to conceptualize this in terms of gallons.

I also didn't know the size of a barrel and would find it useful to get a grasp of this size

Perhaps you could estimate the size of the barrel! It wouldn't be too difficult if you could assume that a barrel was 3ft tall and 2feet in diameter.

I would have thought to make an estimate of what % we make ourselves and then estimate how much we use.

I didn't follow how we accounted for the percentage of use of oil not included in our estimates.

You could also estimate how much the largest exportes of oil produce each day and estimate that the US is one of their biggest buyers...possibly even 50% of their exports go to the US and multiply out that by 365 days.

hmmmm....do you know all the oil producing countries? And what percentage of their exports go to the US? I agree that it's another way of doing, but a much harder one.

I agree too. The way done in the reading seems to allow us to make better (more accurate) estimations.

on two quantities, the diffraction angle and the wavelength of visible light, so its tree has those two nodes as children. All three trees combine into a larger tree that represents the entire analysis:

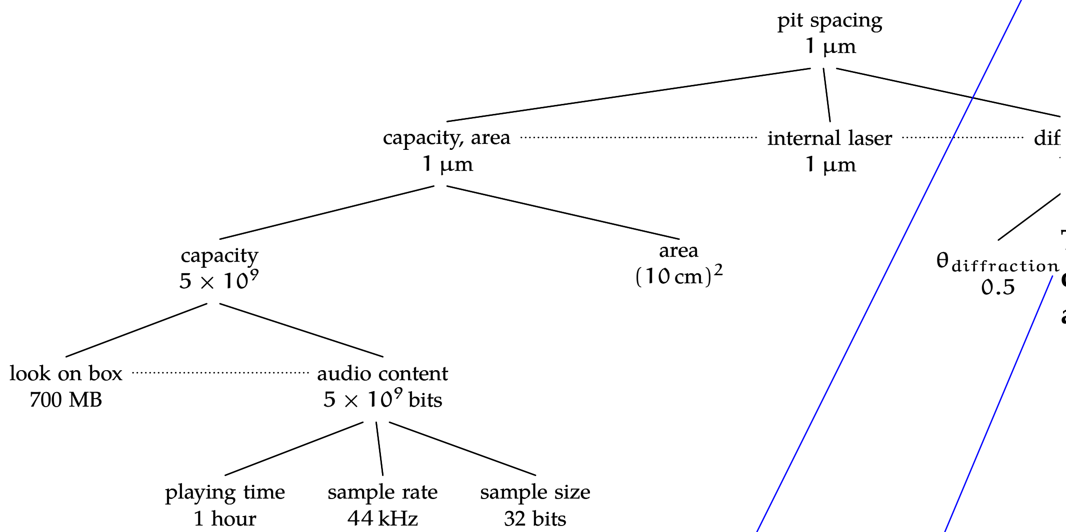

This tree summarizes the whole analysis of Section 1.1 and Section 1.2 – in one figure. The compact representation make it possible to grasp the analysis in one glance. It makes the whole analysis easier to understand, evaluate, and perhaps improve.

1.4 Example 2: Oil imports

For practice, here is a divide-and-conquer estimate using trees throughout:

▶ *How much oil does the United States import (in barrels per year)?*

One method is to subdivide the problem into three quantities:

- estimate how much oil is used every year by cars;
- increase the estimate to account for non-automotive uses; and
- decrease the estimate to account for oil produced in the United States.

Here is the corresponding tree:

Does the US produce enough oil per year to actually necessitate this? My impression was that the US imports a lot of oil and only makes a little bit, so it seems this may be unnecessary?

Having worked for a natural gas company, I have a bit of insight into this problem. The US actually produces 50% of the oil that it consumes, and imports the rest (with Canada being the top importer).

Wow that's really surprising - we produce half the oil that we consume? I guess this will come up later on in the reading...

This probably would not have occurred to me, given that a lot of what we hear about is our heavy dependence on foreign oil, although in the end we would only be off by around a factor of two

Without considering this he might as well have asked how much oil the US uses in a year because we'd be assuming its all imported.

I guess in addition to not trusting the back of CD cases or coffee bags, we certainly shouldn't trust the news or politicians.

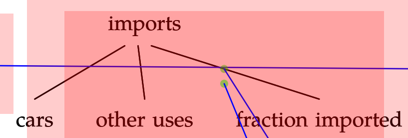

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

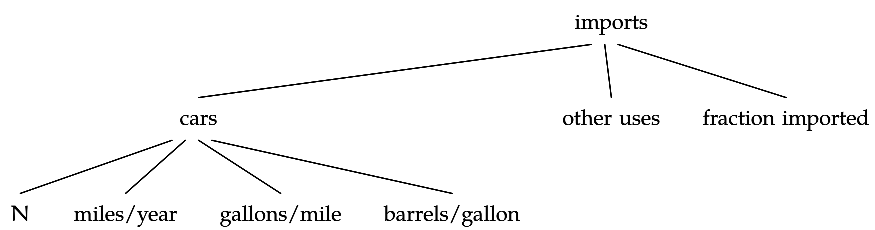

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

What ever happened to swallowing the biggest frog first? (That is, doing the hardest task first.) Is there a particular reason to do the easy ones first? I would personally attempt to do the hard part first, because if that doesn't work out, the easy parts are a waste of time and we need a new estimation method.

I think it's a matter of time. Why spend a lot of time working on one hard thing when you can do many easier things instead?

Perhaps it's just a confidence builder...

I agree. I would think that expanding cars to see what is needed (what children it will have) to calculate it would perhaps be better since it would give the whole picture first.

Perhaps the intention of this statement was that doing smaller calculations will help us see the bigger picture. This is just phrased so that it seems like we are avoiding the most difficult calculation until we have to get to it.

Could these categories be better organized such that we can easily estimate the major consumers (ie transportation and energy production, for instance) and then ignore the lower contributing factors?

I think for your textbook the current tree is great though because it makes for a good illustration of your points - I'm just wondering for solving problems, which is preferable.

I really should read ahead before I post. The next point answers my question about the two opposing tendencies.

One thing that confused me at first was that I thought the branches were combined through addition/subtraction, but they also used multiplication. I would have tried to find values for each of the branches instead of finding one and figuring out how the other related to it (independent versus dependent branches). It seems this way that you could check using redundancy; if I calculated a value for the non-car branches and also found the multiples, I could confirm my estimates.

I somewhat agree. I think that for a simple tree like this, it is easy to remember what is added and what is multiplied but in a complex tree, there might be some value in having a convention that indicates whether values are additive or multiplicative.

Agreed. If it weren't for the other comments already here, I wouldn't have been able to correlate this graph with the math it implies. "Fraction imported" somehow doesn't feel right being on the same branch level.

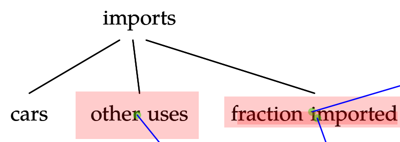

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

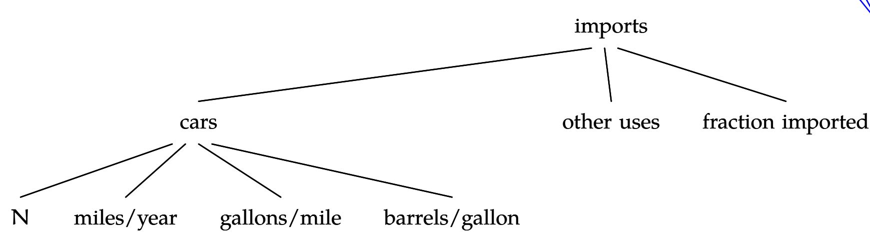

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

Why do we say "fraction" imported instead of "quantity" imported? I would think a fraction would be harder to calculate.

I'm also confused by this since the other 2 values will be measured as quantities.

It looks like maybe this "fraction imported" node represents the "decrease the estimate to account for oil produced in the US," so we will multiply our final answer from the other branches by this fraction. This seems like the case because the other two branches represent the other two subdivisions of the problem that is outlined on the bottom of the previous page.

the point of this is to estimate the quantity imported...having that also be one of the branches wouldn't make very much since. Also, it's quite simple to guess at how much of our oil usage is imported as a fraction.

I'm agreeing with the first two respondents here - I think I understand what the meaning of this section is, but it would be better if all nodes located at the same level of a tree had the same units to reduce confusion.

What does fraction imported mean? How is this different from imports (the top branch)

I would like to hear a little more about the specifics of other uses.

I think it's pretty clear. anything else that requires oil, but isn't cars...

Didn't the tree skip a step? We are trying to calculate the amount of oil used by the us then subtract us production in order to get imports, right? Shouldn't "fraction imported" really be the fraction of oil produced in the us? I thought this graph corresponded with the above three bullets... And if production is what this represents, how would you graphically represent subtraction in a tree?

If you said you were going to begin with the second and third quantities, why is the first quantity the second one discussed?

While I agree that we could look at it this way, i feel like another reasonable stance would just be to guess that cars represent about half of our consumption.

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

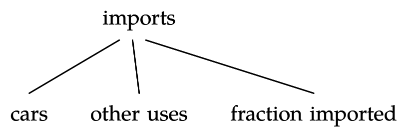

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

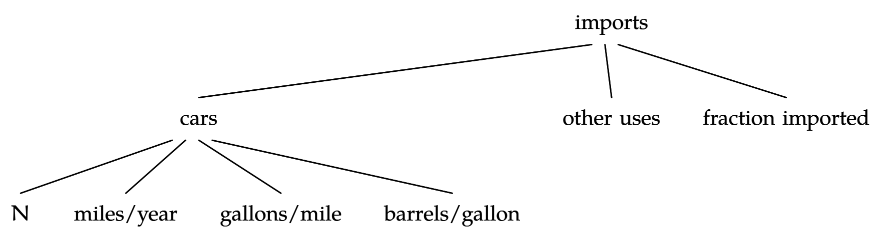

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

I would not say that the automotive uses are virtually zero- it is true that cars use much less gas than planes, but there are also considerably more cars on the road every day than planes. This would lead me to only assume tendency (2) and so I would have "cars" and "other uses" marked as intelligent redundancy.

The wording of this is confusing to me.

I also have trouble with the wording of this. How could the fraction be pushed to zero? What does this mean exactly.

Basically it means this: they are trying to guess the fraction of oil used for cars. The two extremes are that most oil goes towards non-automotive uses, or that most oil goes towards automotive uses. He argues that since both are equally believable, he'll guess that half the oil goes to automotive uses - splitting the difference.

I like the concept used here. I agree that is a very plausible to use 1/2 for non-automotive usage, and it just seems like common sense that would work. I'm not sure if it was included in heating and cooling, but (and this is much smaller these days) power plant usage may also be worth a mention.

I think this should be unimportant instead if you assume a fraction of zero.

I do not see this as a good reason for why they are using 1/2.

me neither, cars use oil, so why would an argument go for a zero value? I understand that you wished to introduce a way of thinking about how much us oil consumption was used by cars, but this doesn't make sense.

it's not that it would be zero, but some might argue that (in the grand scheme of things) the amount of oil consumed by cars is negligible compared to the total consumption. it's like when you're looking at 10^40 + 10 ... the 10 really doesn't count for much at that point so we can ignore it, but it's still not 0.................as for a good reason for it to be 1/2, read the comment 'i would not say that the automotive uses are virtually..."

What's unity again? Is it infinity?

unity is totality

Or in other words unity is 1, since we generally encounter it in relation to fractions.

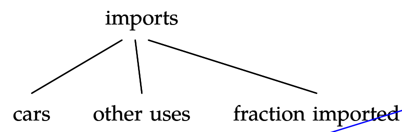

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

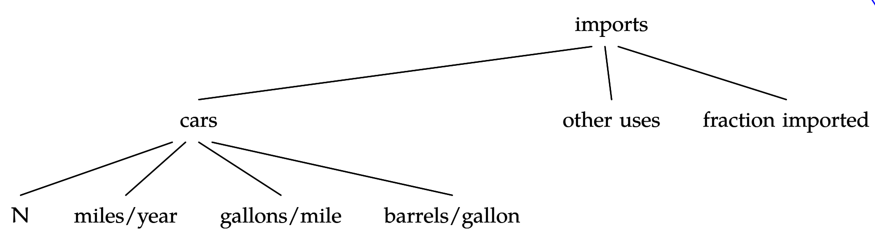

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

I don't know, I thought planes consume a considerable amount of fuel. I'm sure our approximation will somehow account for this but I still cannot mentally dismiss plane fuel use so easily

I'm working on consulting interviews right now and a recent practice case I did estimated that plans use 10 gallons/mile - given how much weight they're carrying and how heavy they are, that sounds quite reasonable to me. Given that, it does seem that planes would take up an extraordinary amount of fuel.

But you have to consider that there are orders of magnitude fewer planes moving at any given moment than automotive vehicles. I'd even hazard to guess that the fuel used by the trucking industry alone in a given unit of time would be at most an order of magnitude smaller in comparison to that used by planes (assuming 10 gal/mile is reasonable).

Okay, this may just be for me, but I have to reword this to make it clearer and to explicitly define the relationship show in the tree structure—

(1) the idea that the non-automotive uses are important enough to draw the fraction of cars/total toward zero; (2) the idea that the automotive uses are important enough to push cars/total toward unity.

I also found this confusing. I would suggest making the statement sound more hypothetical.

(1) The idea that the non-automotive uses might be so important that they overwhelm the contribution from automotive uses; (2) the idea that automotive uses are so important as to overwhelm the non-automotive uses.

And maybe then talk about the effect on the fractions after you have defined what these two cases are.

So are we forced to pick one? Which is more important? non automative uses or car uses....When you say 1/2 what does that mean? Are you considering half considering both? Or picking one? I am bit confused by that comment.

What he is saying here is that the non-automotive uses are roughly equal to the automotive uses, so both uses each take 1/2 of the oil.

He is taking the intermediate stance that "other uses" account for as much oil as cars, so he just doubles the estimate for oil use from cars to get oil from other uses+oil from cars.

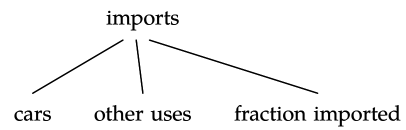

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

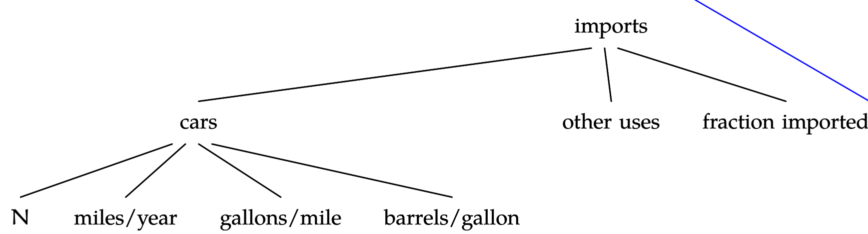

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

Overall good example. After reading the comments, it seems to me that you may want to explain why you can be so loose with the percentages you guess.

It seems like a lot of people are having trouble accepting the 1/2 number. I think if you wanted to, you could relate it back to everyday life (i.e. think about how much we travel in cars, trains, trucks, planes, etc.) and break it down by how much fuel/person each mode of transportation uses. Then I suppose the oil used for manufacturing hydrocarbon-rich products would be much more difficult to estimate.

Are we allowed to trust our gut to that extent? Reasoning that seems reasonable to ourselves?

And had there somehow been 3 different areas, would it have been 1/3

I think this is reasonable, it would be tough to come up with a better fraction.

I agree that I wouldn't have been able to come up with a better fraction, but the decision to go with 1/2 seems arbitrary. Is there another way we could come up with this fraction?

Does that mean we could guess any fraction and continue with the problem with that assumption? I feel like this is an area that would get me stuck, because I'd be trying to figure out the details rather than making a guess. How do we know when to guess and when to find a more educated approximation?

Is the one-half estimate just a very rough guess or was there more too it? How much freedom can we take when approximating numbers this important.

Why? I understand that a lot of the listed things are larger than cars, but there are a LOT of cars on the road. Given that today we assumed the professors were an insignificant portion of the total staff in the school, why would it be unfair to assume that the sheer number of cars used on a daily basis (and their old needs) do not outweight the use of oil in trucks and planes by over 1/2?

because we are making 2 assumptions here: either the non-automotive uses are important (hence we can ignore the use by cars) OR the automotive uses are important (and hence we ignore the non-automotive), since 2 assumptions are equally plausible, we weight them equally and thus 1/2

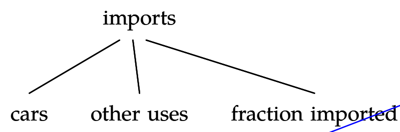

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

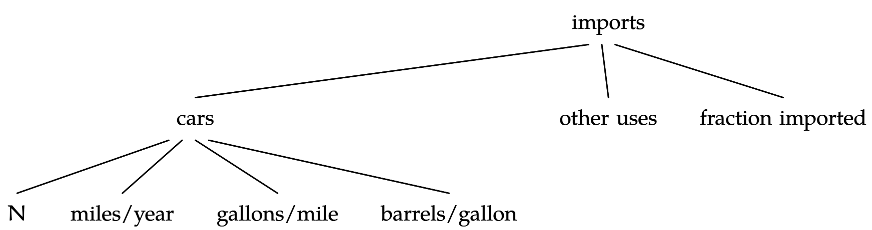

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

I would agree with this estimate, I think that the non-automobile uses are very important and although these are less things, they each use more oil

seems like the accuracy of any kind of estimate here is very likely to be off by alot

This sounds reasonable but I would assume that cars use more then half the oil in this country being that there are so many. I feel as thought 1/3 would be a better estimate for non-automotive uses.

how is this estimation reasonable?

I get the 1/2 – but I don't get why we double the oil consumed by cars.

I don't get it either...what does non-automotive uses have to do with oil consumed by cars? They're sort of different by definition...

I believe the idea is that since we think it's equally likely that cars make up a vast majority and vast minority of oil consumption, we have to guess that they account for half (the middle between our two estimates)

I see that the wording in the text is not very clear. What I meant: The imported oil is used by cars and by others (trains, generators, trucks, heating, ...). The factor of 2 is my guess for the factor to increase the car estimate to arrive at the total estimate.

As with others, it makes sense that we conclude that the car fraction for oil consumption is 0.5, but when you say double it, doesn't that mean you end up with 1? As in all the oil consumed is consumed by cars? It seems like this last sentence has really confused people.

It was confusing. What I meant is that cars consume 0.5 of all oil, and imported oil is 0.5 of all oil, so car consumption is roughly the same as oil imports.

I am also confused as to how this a reasonable estimation. I'm assuming that it turns out to be pretty close but I don't understand how just because I don't have a good feeling for two potential options that I can just assume it to be half one and half the other? what is the justification for this?

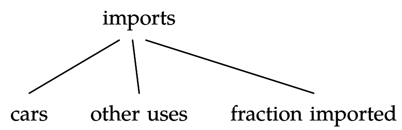

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

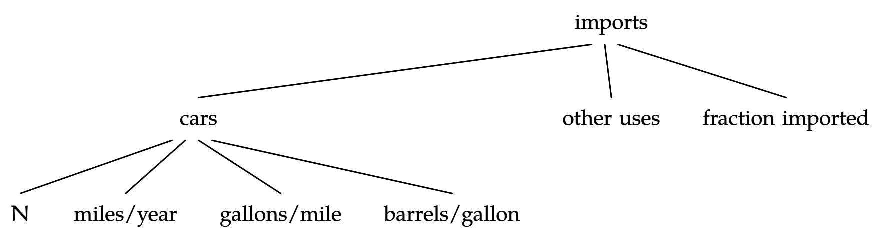

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

this is confusing. Exactly what estimate are we doubling? and why?

This phrase kind of makes me feel like imports are somehow consuming oil, which is of course exactly the opposite of what you mean.

Perhaps, 'imports make up a large fraction...'.

I'm still not so sure that this is a valid explanation. I think we read about a lot of things in the press that get exaggerated.

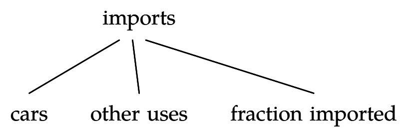

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

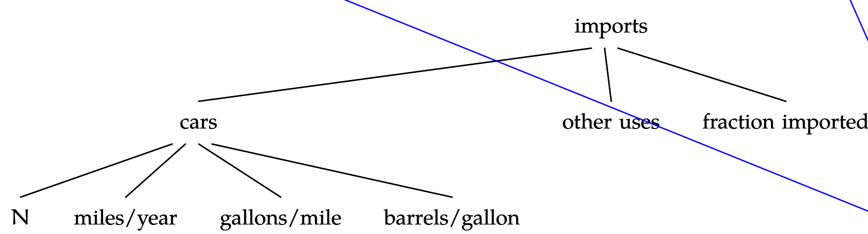

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

i agree that imports are a "large fraction of total consumption." thus, if i were to estimate this, i would guess a much larger fraction of imported oil, maybe 2/3rds to 80% or so. why did you choose 50%?

I agree

I also agree. I thought imported oil would be closer to something like 90%–that seems like high enough a number to have it on the front page of the newspaper every day.

I disagree. 50% is a very significant percentage. Imagine if you were all of a sudden paid half as much at your job, or alternatively only half the people could be kept on staff. If 50% of the company was getting laid off, would you be worried?

I agree that 50% is significant. We are already assuming that 50% is used for cars, and the other 50% is used for non-automotive processes. Thus, if we lost all of our imports, it would be the same as losing the car branch or non-automotive branch.

I also agree that 50% is very significant; I think 90% is way too high, if that were the case, there would be so much more U.S. control and involvement in international countries than even now.

So I actually looked the percentage up out of curiosity and as of 2007 66.19% is imported.

%50 seemed like a low number at first but the above comments convinced me. %50 of the oil industry is enough to receive so much attention in the press. It's strange that domestic oil doesn't receive much attention.

I'm just not sure it matters all that much, at some point. Taking 50% of a quantity is, order of magnitude, not that different from 90%. Especially when there are so many other rough estimates being made.

So if the person lacks common knowledge, how would he go about making estimates? Literally just ball park it?

Yes. Any person must have some sense of numbers in the world, and a gut-aided guess it better than none.

agreed

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

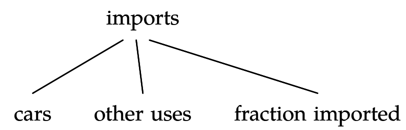

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

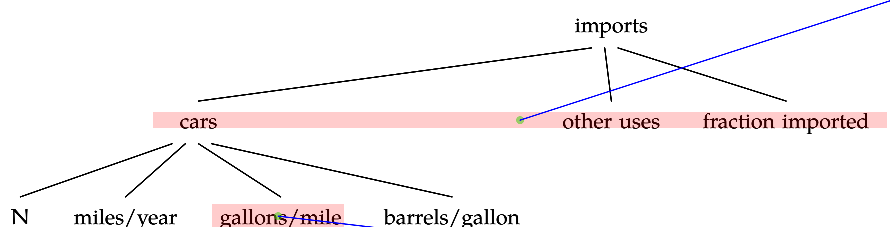

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

I think this is pretty reasonable too. If I had to guess, I'd say more than half is imported.

On the first page of this reading, one of our classmates who worked at a natural gas company also said about 50% is imported, so this agrees. Hooray!

So at this point we think that half of our oil consumed is imported yes? So if we also think that cars are half our consumed oil. Can we simply say oil consumed by cars=imported oil?

Interesting point. It seems reasonable enough considering the way we're estimating, and this would make the estimating a lot easier, since would would have to double and halve.

those estimations introduce error in specific ways to our final value, so we can't simply ignore them...meaning even though they cancel each other out, they are still in the tree for reference.

That makes sense. Though I was going to comment that 1/2 seems a bit low; maybe 2/3 or 3/4 seems a bit more accurate, but I may be way off.

Though, what you said makes sense, and that tree could be 'cut down'.

How did we come up with 50%? That seems like a pretty large number.

do these all have to be values that can be multiplied together straight across the tree? can you set up a tree that needs different operators between the branches?

No, two of these (cars, other uses) could be added. The branches (quantities) can be linked by any operation.

This is why in the previous section people thought it would be a good idea to somehow indicate what the connection is between each leaf and branch in the tree, operation-wise

I agree. Otherwise I'm afraid we'll get overlapping estimates.

With the very wide range of MPG ratings, especially with the difficulty involved in estimating the value for hybrid cars, how difficult is it to arrive at this number?

I'd guess that it's safe to go with a slightly lower MPG rating, since a lot of cars on the road are not hybrids. We know that a sort of typical MPG rating is 20, and nothing is going to be too dramatically higher or lower than that.

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

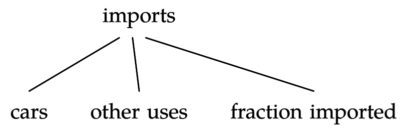

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

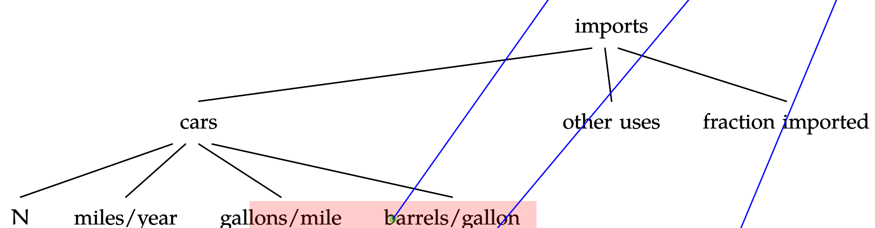

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

Shouldn't these two fractions be reversed? As in, miles/gallon and gallons/barrel?

That's how you wrote them in the text above, so it seems strange to reverse them here. Also, maybe replace 'N' with 'number of cars' for clarity.

Ah, I see later that you use them in the written order to make the units work. Maybe you should specify that somewhere?

This seems like a difficult estimation problem in itself. How would we even tackle this?

I don't know about this. Every person has a car? Why not every 2 people? I think that's a safer guess.

Even if it is correct that every 2 people has a car, your answer would only be off by a factor of 2. I think the important thing is that your answer is still within the correct order of magnitude.

What about commercial vehicles and other vehicles that aren't being used for personal use? Vehicles such as Semis and company owned trucks consume much more oil than personal autos. I feel like this easily drives this ratio much higher, at least 1:1

What about commercial vehicles and other vehicles that aren't being used for personal use? Vehicles such as Semis and company owned trucks consume much more oil than personal autos. I feel like this easily drives this ratio much higher, at least 1:1

commercial vehicles are an interesting thought...it doesn't just increase the person/vehicle ratio, it raises the miles /year (a lot), and raises the gas/mile (a whole lot)...

maybe this could be part of other uses

I actually asked myself the same exact question. But from daily life you can roughly tell that roughly, there is 1 car per 1or2 people. Even if you are off by a factor of 2, it makes sense to think that you are still within the right order of magnitude.

Is this to compensate for the number of people that have multiple cars?

what would it matter if they aren't consuming oil?

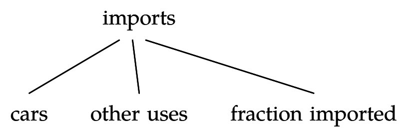

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

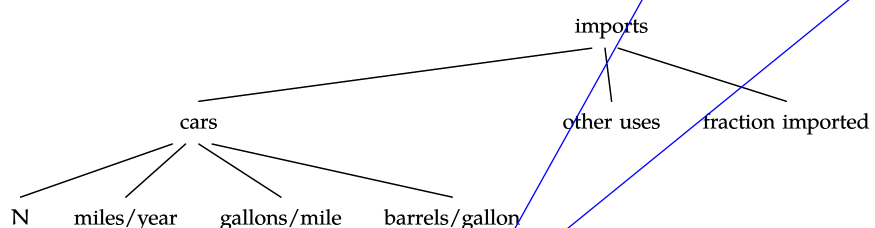

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

this is a bit of an exaggeration. I would have said that the average family size is around 4 people, and each family has on average 2 cars. That seems more plausible to me.

agreed, but this only changes things to, say, 1.5*10^8 people, which is only a difference of 10^.5, which is probably negligible for an approximation

It is also necessary to consider the number of cars owned by companies and other 'commercial vehicles'. Thus, I don't believe it is too much of an exaggeration, if at all.

I don't think that the number is that far off, but the explanation of how it was estimated is a little strange. The average person would most like take the families + businesses etc route. For me, just saying "even babies have cars" leave out a large part of the logic behind the number, making it seem like more of a lucky guess than an educated one.

I'm pretty sure it's a joke about how wasteful the US is... that said, the approximation does happen to account for commercial vehicles and personal ones.

Id guess 1/2 that

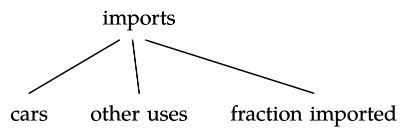

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

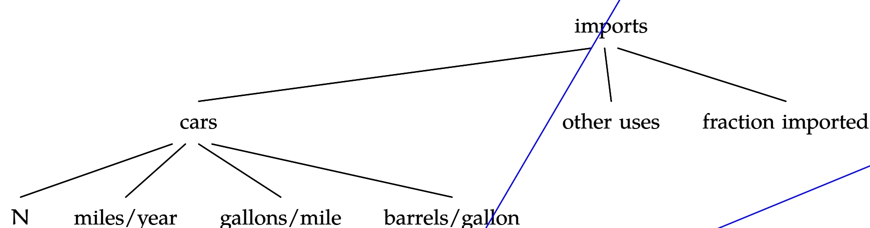

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

But this, as mentioned would include everyone under 16 who can't drive, plus the very old who live in nursing homes. It seems like an overestimate to me!

But then you have people lay Jay Leno with 100 cars. And lots of people have different cars for different uses–say a truck and a car, or a summer car and a winter car. It probably balances out somewhere

But then again, we're trying to get at how much oil is consumed, so it shouldn't matter that Leno has 100 cars since he can only drive one at a time. Therefore, it does seem that including everyone might be an overestimate.

I agree with the fact that this number seems entirely overestimated. Even celebrities with 100 cars are rare- most people have 1 car to their name and then there are some families who do not have 1 car to share so it seems as if estimating 1 car per person (including young children and elders) seems unreasonable and will provide for an overestimate. Also, the point that only 1 car can be driven at a time is a good one to consider- even people with more than 1 car can only drive 1 at a time. I think the estimate should be more like 2.4 million (by subtracting 60 million for children under 16 and elders over 75)

Even with this, there are people that work a rather large district as either a sales rep, technician, etc. that will drive way more than whatever annual mileage estimate we make. I have a few friends that drive all around the Southwest (to the tune of almost 100,000 miles per year) who, by the mileage estimate presented, would be the equivalent of themselves and 9 non-drivers.

I would think that this would be lower, like maybe half but it might be counter acted by the commercial use of cars

Wouldn't it be more fitting to approximate the number of drivers as opposed to the number of vehicles. An individual with two cars in this example is predicted to drive twice as much which I don't believe would be the case necessarily.

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

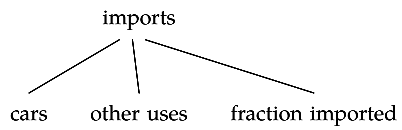

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

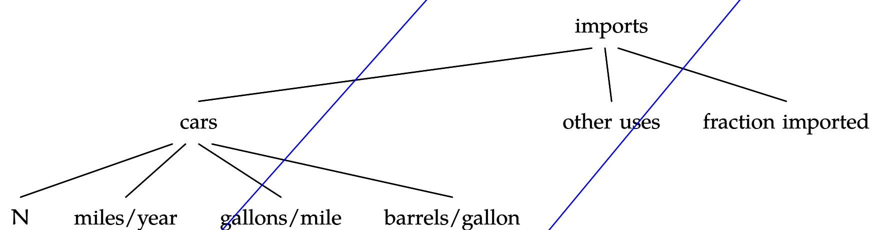

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

I'm not sure if anyone else could have assumed this, but it seems like a difficult number to instantly assume. I think an easier way to get an estimate would be to guess how many miles a car goes a day, roughly 40 would be my guess... and from there get the number 15000

15000 per car seems high to me, especially if we're counting babies' cars. That implies 15000 miles per person (no carpooling) per year? Maybe half of people aren't driving significantly in a year (too young, live in cities), so that's more like 80 miles a day per driving person. Is it reasonable that every person that drives sits in their car for nearly 2 hours a day?

I can't imagine 3*10ˆ8 cars being driven 10K miles/ year. I'd say each family having one car that goes 15K is a better estimate.

I disagree. I actually like this estimate. It seems to me that most families own 2 if not 3 cars.

It's also worthwhile to note that most new car warranties are 100,000 miles or 10 years.

my family (2 cars) travel 120 miles/weekday to commute to work (1 hr & 1 half hour commute)...that's 15K miles/year for 2 cars. Both of my parents have fairly standard commuting times for the area that I live in...I'd say that the estimate can't bee too far off, especially considering that those numbers don't take into account the fact that my dad works 6 days a week and they go places other than work...

How do we assume 15,000? just try to use personal experince?

Experiential knowledge works, but you can also do a quick estimate: 30 miles per day, 7 days per week, 50 weeks per year = 10000

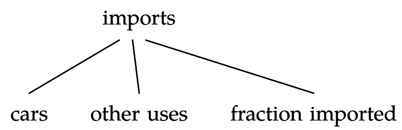

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

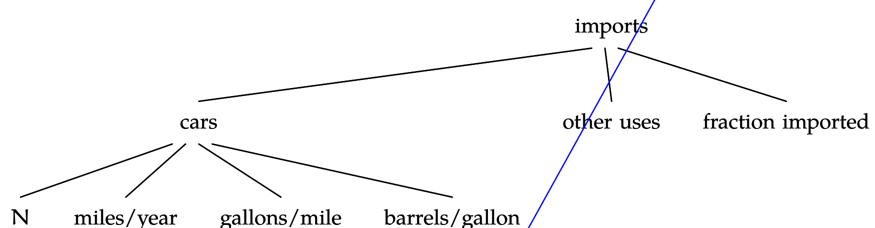

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

If I guessed 10000 from the beginning, then compensated for that numbers, that may lead to an even greater deviation. But I guess it all depends on an initial guess. One guess may trickle down and lead to several other possibly erroneous estimates.

I'm also a little hazy on this. How do you decide if your guess is accurate enough to afford adjusting it to compensate for other numbers?

I feel like in this case given that the number of people is an overestimate, lowering the annual miles per car should sort mitigate the errors because they are errors that cancel... but I'm not sure if that was the justification intended for it.

I still feel that I would've begun my guess for miles per year at 10,000 since that is the usual value used for car warranties as advertised by dealers - 5 yr/50,000 mile, etc.

I think making that approximately is fine - we're still on the same order of magnitude, and only off by 1.5x, which for the sake of the estimation is pretty insignificant.

Also, if you think your initial guess is high why not just adjust it instead of potentially introducing more errors on other variables? I know that you can compensate for over estimates in one area by underestimating in others but I dont really understand how this won't lead to compounding or potentially multiplying your errors.

The trick is you don't know if your overestimating or underestimating. At the end of the day, an approximation is an approximation. You can only get so close to the correct answer

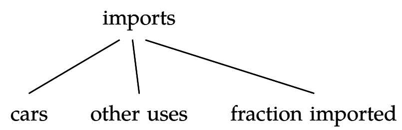

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

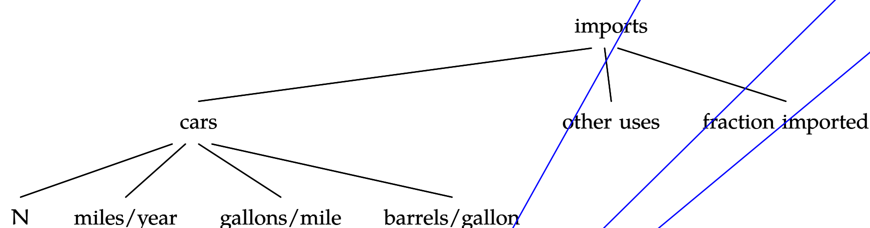

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

These seem like good estimates. Most people only own 1 car but some own more and as far as orders of magnitudes go, this would make sense.

I think that estimate is a bit on the high side. I know in the wealthy suburbs, families usually have at least 2-3 cars (though some have 4-5); however, these familes usually have 4-5 members and this is the upper end of middle class. In the poorer parts of Chicago where I used to live, many familes (larger too, given economic conditions) don't even have cars at all. In fact, in cities, even many wealthy people never bother to get cars because it's far too inconvenient. Although rural America >>>> urban America in size, population densities make them about even. So if I were to guess, I'd say maybe 1x10ˆ8.

My estimate does neglect non-consumer vehicles though, like mass transit, delivery, and work vehicles. So if you x2-3, that does you get around 3x10ˆ8.

I agree that the estimate is a bit high, but definitely not more than an order of magnitude so for this class it seems fine.

Nevermind, I guess this makes up for the overestimate.

There isn't one number here that I completely disagree with, but I feel as though none of them are really known other than the 3*10ˆ9 people in the US. I'm also still not completely comfortable with trying to manually account for possible errors by making more possible errors.

Accounting for errors with other errors seems to be one of the key ideas in this class, so long as you know in which direction you're erring in.

I agree. accounting for errors by over or underestimating other values seems a little shaky, but sometimes that's the best you can do. if you know you're overestimating something, then you can try to underestimate something else to compensate. obviously, this won't be exact, but after all, this class is about estimating when we don't know how to solve a problem for an exact answer

I found this paragraph a tiny bit dense, just because there are a lot of facts thrown out, a lot of which I didn't know.

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

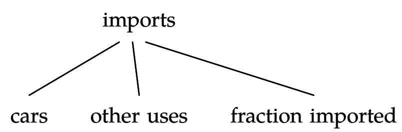

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

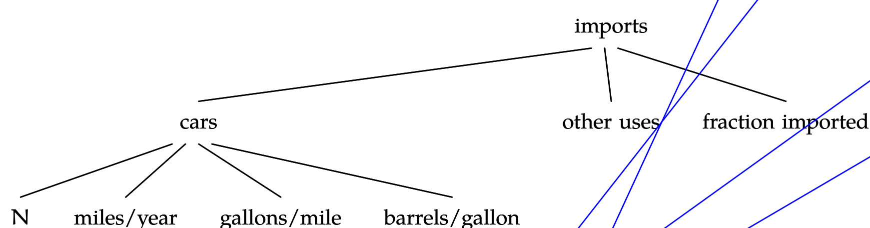

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

i know a lot of cars who claim that 25 mpg is "energy efficient" or "above the standard." i would have guessed this number to be lower, especially given america's tendencies towards suv's.

I agree...most cars get usually 18 mpg...its only a new fad that cars are more energy efficient...and most of the cars in the US are used.

Also, I think a lot of driving is done in the city versus freeway....thus the mpg would most likely not crack 20

While I agree this is high, it is not incredibly far off. If you said 20 your off by 20%.

This seems a bit high too. Typical listed mileage might be 25/gallon, but the actual mileage has to be lower.

I agree... while we see this numbers like this all over the place in auto ads, these are all hwy millage and not nearly what the average person gets a day. What about guessing the tanks of gas used a week and the average tank size?

that would be _way_ more complicated that you'd want to think about. i'd say 25 is a good estimate...of my family's 4 cars we get 15, 12, 48(hybd suv), & 20...and we drive 3 old junkers.

this answers the question above.

This appears to be a legitimate way to calculate how much a barrel holds...however, the cost that we pay for a gallon of gas is greater than the actual price of a gallon. I don't know how much the price is raised, but it could make a barrel's capacity greater than 40

Good point. Is there any way to actually calculate the premium we pay? It seems that it is a large factor.

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

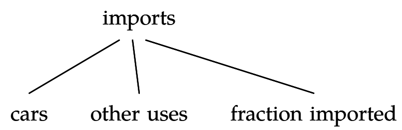

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

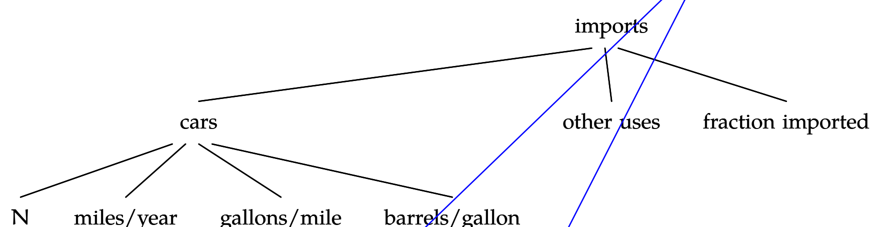

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about $100, and a gallon of gasoline costs about $2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

If the idea of a barrel is similar to that of a trash barrel... the typical drum liners, or trash bags we use in our fraternity are about 50 gallons...making a barrel slightly smaller, which agrees with this estimate of 40.

Yeah, that's more or less how I would have approached this one. It's probably more accurate than my barrel/gas-cost guessing abilities.

I would have had to look up the coast of a barrel, but the other analysis was where I was headed when I first read the question. How come you didn't break up gallons/barrel into another tree - I would've done it that way.

I personally have no intuition about the size of a barrel and could have easily guessed $1000 for its cost, putting my estimate off by a factor of 10. How do we calculate a reasonable estimate if we are unfamiliar with the units requested?

Yeah I had this same problem.

Oil hasn't traded at $100/barrel in a while - I think $75 would be a much better estimate given the past few years. However, this is not even an order of magnitude off, so I guess it doesn't matter.

is a barrel the processed oil that we would buy from the pump, or does it have to undergo further processing, raising the price?

I think where some of the confusion comes in here is that there are so many different factors that the estimate has not accounted for, which is ok, since it is just an estimate. For example a barrel of Oil, 42 gallons, only produces about 20 gallons of gasoline after refining, And the fact that the price of a barrel of Oil and the price of a gallon of gasoline can vary so much because of the time delay, refining bottlenecks, and distribution that the fact that the estimates they made turned out close to being correct is either luck, or just good use of averages.

how do we know that a barrel is $100? I thought that they were more than that

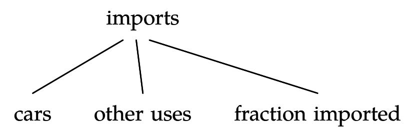

The first quantity requires the longest analysis, so begin with the second and third quantities. Other than for cars, oil is used for other modes of transport (trucks, trains, and planes); for heating and cooling; and for manufacturing hydrocarbon-rich products (fertilizer, plastics, pesticides). To guess the fraction of oil used by cars, there are two opposing tendencies: (1) the idea that the non-automotive uses are so important, pushing the fraction toward zero; (2) the idea that the automotive uses are so important, pushing the fraction toward unity. Both ideas seem equally plausible to me; therefore, I guess that the fraction is roughly one-half; and, to account for non-automotive uses, I will double the estimate of oil consumed by cars.

Imports are a large fraction of total consumption, otherwise we would not read so much in the popular press about oil production in other countries, and about our growing dependence on imported oil. Perhaps one-half of the oil usage is imported oil. So I need to halve the total use to find the imports.

The third leaf, cars, is too complex to guess a number immediately. So divide and conquer. One subdivision is into number of cars, miles driven by each car, miles per gallon, and gallons per barrel:

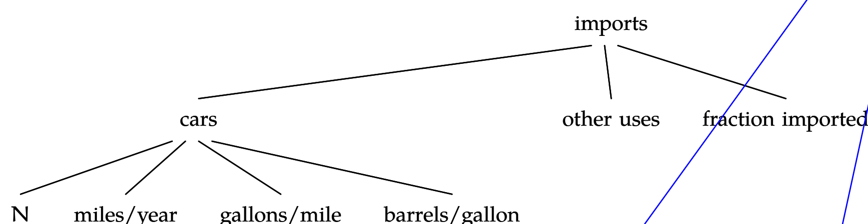

Now guess values for the unnumbered leaves. There are 3×10^8 people in the United States, and it seems as if even babies own cars. As a guess, then, the number of cars is $N \sim 3 \times 10^8$. The annual miles per car is maybe 15,000. But the N is maybe a bit large, so let's lower the annual miles estimate to 10,000, which has the additional merit of being easier to handle. A typical mileage would be 25 miles per gallon. Then comes the tricky part: How large is a barrel? One method to estimate it is that a barrel costs about \$100, and a gallon of gasoline costs about \$2.50, so a barrel is roughly 40 gallons. The tree with numbers is:

Interesting fact...the cheapest country with gas prices is Venezuela (where I'm from)....gas is about 0.29 cents a gallon...cheaper than a liter of water.

Accounting for the different classes of gas (regular, medium, superior), the price seems like it should be more like \$3.00 and even that might be an underestimate these days.

I think this is more of a wholesale cost rather than consumer cost.

Or perhaps this is just off because of the crazy fluctuations in gas prices in the last few years, but even so the difference shouldn't change the final answer drastically.

Don't the gas companies buy the barrels for less than they sell them for? (Hence making a profit) So estimating this would be an underestimate of the actual amount in a barrel?

Is this a credible price estimate after accounting for fluctuations over time?

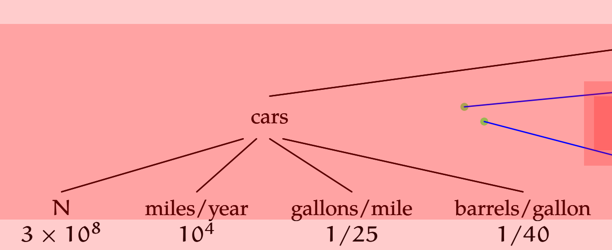

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^{8}\ \text{cars} \times \frac{10^{4}\ \text{miles}}{1\ \text{car–year}} \times \frac{1\ \text{gallon}}{25\ \text{miles}} \times \frac{1\ \text{barrel}}{40\ \text{gallons}} \sim 3 \times 10^{9}\ \text{barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^{9}\ \text{barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^{9} barrels/year – only 25 higher than the estimate!

> **Problem 1.5 Midpoints**
> The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

Comments on page 3

Basically, we are going to take the figure from oil and then adjust it. I think the notation is confusing since it does not take adjustment into account

I made a comment earlier about using the gallons consumed by cars and then guessing the fraction that we import. It seems like this method is a little different that originally described

does the proper notation of these trees include numbers or no?

I realize this doubling is discussed in the text, but wouldn't this be better named fraction used by cars? (50%, so we multiply by the reciprocal) Or it could be specified that the branches are equal (cars = other uses). That would make it easier if our redundant method involved calculating the other uses and assuming that cars made up the other half of the total.

Also, if I were to just look at the graph, the way it is worded makes it seem like there are 2 other uses, not that the other uses account for 50% of the total amount used.

so long as you remember that you're starting with a figure for oil used by cars and multiplying it by adjustment factors on the same level leaf. The naming is just to help you remember why that factor is there.

I agree that the notation on the other uses and fraction imported leaves is a little confusing. Indicating operations between the leaves or simply calling them "other use adjustment factor" and "fraction imported adjustment factor" would clarify the units on the figures in these leaves.

Hey, so if we all stopped driving, we wouldn't need foreign oil. Secret to world peace: bicycles.

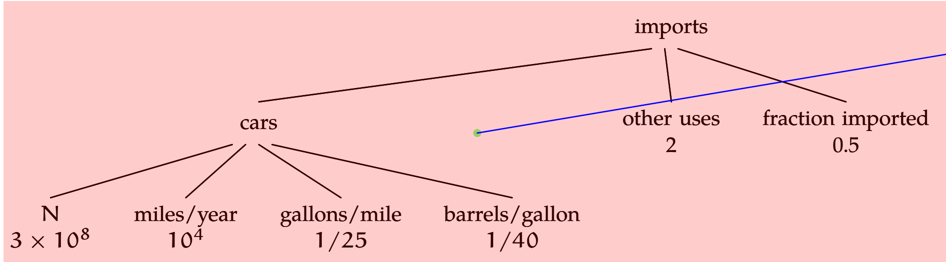

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3 \times 10^9 \text{ barrels/year.}$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year.}$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

This math seems to work out too easily, everything cancels to give 3x10^8 x 10.

There's nothing wrong with that. This is an approximation class and I guess there's an 'art' to it. When you don't have a calculator, it's best to try to break things down into easy numbers w/o losing too much accuracy and the easier you make it for yourself, the faster you can make an estimate.

I agree with the comment above. Why do you feel the math has to be complicated or messy to be right? As we see a few lines down, the estimate derived "too easily" is very close to the actual value, and it didn't take us nearly as much effort as calculating it exactly. While I do agree some things require precision (i.e., anything where lives are at stake), I don't think we should rule out estimation for it being "too easy."

Like in the previous section, I think it would help if the operations were incorporated into the tree. If the paragraphs/text were taken out and one only saw this tree, having the operations would make this diagram much clearer and one can simply use this tree to solve the problem.

$3x10^11 of oil each year. I thought it would be more.

That's a pretty awesome estimate. For someone who knows about energy, but didn't have any statistics, I would have gone with these numbers: 2*10^8 cars (so 1*10^8) (One car per household, 2 people per household, *2 to include trucks and public transport) avg drive 40 miles a day, 1200 a month, 15000 a year (would have left it at that), would have done down on the mpg; maybe 20 (lots of trucks and older cars), and I would have used 42, but it'd round the same. So I think it's about the same; but that just seemed very elegant.

essentially dimensional analysis is equivalent to the "divide and conquer" method?

huh?

The "car" leaf determined the barrels/year for cars used by households. "Other uses" leaf would consider the number of cars by businesses and "fraction imported" is a rough estimate of how much oil is imported. Luckily 2*0.5=1.

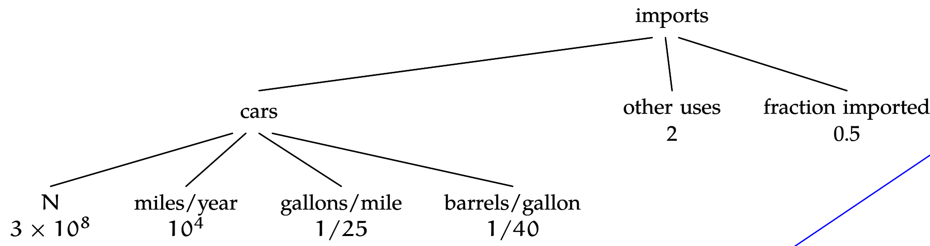

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3\times10^{8}\,\text{cars}\times\frac{10^{4}\,\text{miles}}{1\,\text{car–year}}\times\frac{1\,\text{gallon}}{25\,\text{miles}}\times\frac{1\,\text{barrel}}{40\,\text{gallons}}\sim3\times10^{9}\,\text{barrels/year}.$$

The two adjustment leaves contribute a factor of $2\times0.5=1$, so the import estimate is

$$3\times10^{9}\,\text{barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^{9} barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

I don't understand this math. The cars node contributes 3x10^9 barrels per year, but aren't you supposed to just x2 to get total since we estimated above that cars used up about 50% of oil? How are we supposed to tell from the tree that you are supposed to multiply cars, other uses, and fraction imported?

Just multiplying the cars consumption by two gets us the estimation for all the oil used in the US, but we decided that 50% is made in the US, thus not imported. We are trying to find the total imported, so we only want half of the oil used in the US.

Clearly, it is best to assume half the oil usage is not automobiles but it could be that more is imported then we took into account.

I don't really understand what just happened hear. How do the "adjustment factors affect imports? Does the two represent the denominator of a fraction and if so what is the denominator for cars? 1?

We've found how many barrels of oil cars use each year but that's still only part of the tree. We said that cars are only 1/2 of the oil consumption so we have to multiply our value by 2 and then 1/2 of it is imported so we multiply by a 1/2 as well. These factors cancel out to 1.

It's kinda funny how leaving things out or including them ends up just about the same.

So maybe this is just me being not very proficient at estimating, but sometimes I get the feeling like this text looks at a statistic and works backwards to make its estimation somewhat accurate.

I get that feeling too. But maybe Sanjoy has just done so many estimation problems that he can pick better numbers than we can. Some of these numbers like US population, gallons/mile, etc. are learned through time and the more experience you have with these sorts of problems, the more realistic you'll be when you guess.

I sometimes feel that way too - the estimates are just too close. Even for an experienced estimator, you would expect to be off by an order of magnitude or so every once in a while. He is however, extremely good at estimating, so it's possible he's developed an intuition for what kind of answer "feels" right.

is there a way to calculate uncertainty or is that just a confidence interval that we think? do we make a separate tree of uncertainty intervals

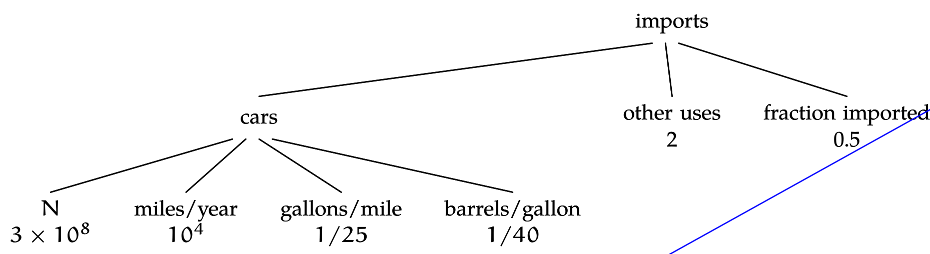

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3\times10^8\,\text{cars}\times\frac{10^4\,\text{miles}}{1\,\text{car–year}}\times\frac{1\,\text{gallon}}{25\,\text{miles}}\times\frac{1\,\text{barrel}}{40\,\text{gallons}}\sim3\times10^9\,\text{barrels/year}.$$

The two adjustment leaves contribute a factor of $2\times0.5=1$, so the import estimate is

3×10^9 barrels/year.

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

Before the oil crisis of 08! People were taking about 200$ a barrel and gas prices at 4.5...which is on average the prices in Europe (they are so much higher than the US)

And, as an interesting note....interest in removing oil from shales in the midwest, which requires using an extraordinary amount of energy to heat the rock up to release the oil (one suggestion, a prof from MIT, was to use nuclear power....)

So our answer is really close, but it would be interesting to know why. Is it possible to, in hindsight, say these particular estimations were bang on, these were too low and these were too high, but here's where the errors canceled? Is that data available?

Well, I'm sure statistics are available on oil use in the different areas (especially manufacturing). I'd guess that a decent amount of error results from our assumption of 50% imports, since I think I've read and heard something closer to 60% (which would give us 3.6 instead of 3, which is almost spot on)

Yeah, thats true. I also had read that roughly 2/3 of the oil was imported. When I did the calculation though I got 4*10^p which wasn't so far off.

I feel like if I did this estimating, I would be very off.

It would be cool here to include the actual values of some of the actual factors we estimated, like the number of cars and amount imported.

I think the whole purpose was to show a basic estimate with accurate results. This shows how we are able to simply the problem without knowing additional facts. Comparing those numbers to actual values of "other uses" and "fraction imported" does not really contribute to the topic.

this should be 25% higher

I'm curious: what fraction of your estimates end up elegantly close to the actual value? Awaiting the mantissas...

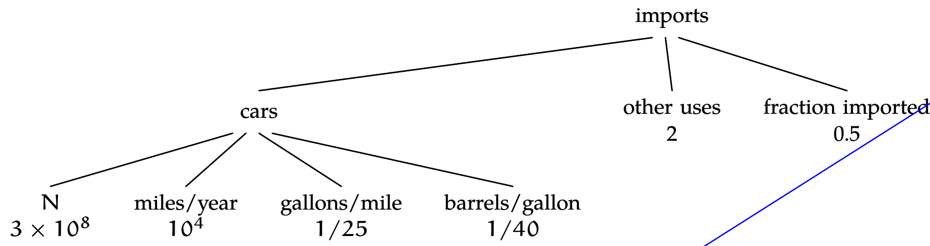

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3 \times 10^9 \text{ barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

typo: should read 25%

How close do we want our estimates to try to be? I guess it's within an order of magnitude, but 25% does seem like a lot compared to some of the stuff we were doing in class, where we got less error (although I guess we also got some guesses that were probably around 25% off or more...)

I feel like with many approximations, the idea is to get within an order of magnitude. In this case, 3 versus 3.7 is pretty darn close.

i agree, this is a close number. it helps when you know what it is before you begin however ;)

It is within an order of magnitude, but still, 25% seems like a pretty big margin of error...

isn't 25% a significant error?

You mean 25% higher, as opposed to 25 higher

but even though it's 25% higher, it's still a good estimate, as we want to look at orders of magnitude. obviously when we multiply by powers of 10, a small error in our estimate becomes amplified. but this is still a good estimate here. we got in the right ballpark: the real value is 3.7 billion, when our estimate was 3 billion (at least we didn't guess something like 3 million)

This seems completely random. We know that we are somewhat working on a log scale because of "few" but why does this belong here? Just to do a proof?

I agree, this "concept question" seems out of the blue and not obviously relevant to the previous example.

This does seem out of place. I think it would have made much more sense if it was placed with the initial discussion of "few".

How is this related to our tree diagrams? I would show that...using log does it have something to do with the quantity of e

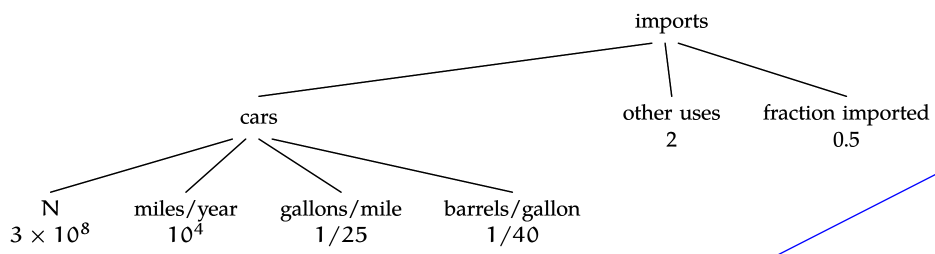

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3\times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3\times 10^9 \text{ barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

this problem was so much simpler! i think it was really good for more practice...it wouldnt have been very good for teaching though–I really like the complexity of the previous one (and the ability to verify different parts).........[there should be line breaks in nb].......I was thinking that it would probably be a good idea to have sidebars in the final text. you see them a lot in math/science texts, include general information about the assumption, like what a pit is or what sampling rates are.

this makes no sense to me

What is the log scale, and how do we calculate the geometric mean?

The geometric mean of 2 numbers n and m is the square root of their product.

GeoMean=sqrt(n*m)

The standard scale would be 1,2,3,4,etc. Using a base of 10, the log scale would be 10^1,10^2,10^3, etc... See http://en.wikipedia.org/wiki/Log_scale for more info. http://en.wikipedia for page on geometric mean (or see the answer in the other comment. I'm unclear how the problem relates the the previous example, it seems somewhat random.

thanks! this helped me too

i dont quite understand how this fits in with the previous section. i dont see how it is an estimation problem either.

Perhaps you can elaborate. In the last section we discovered trees. In this section we used the tree method to estimate something. Do you think this estimate is to obvious?

In the beginning it said that we would subtract for the oil produced in the US, did we decide to neglect it? I didn't see that addressed in the estimation.

That's the "fraction imported" node. I don't think this was clear to very many people. I think a different node name, even simply "fraction of oil produced in US" then multiplying by the 1-(this fraction) might even be better.

'usual scale' seems a little vague since we are talking about geometric mean on log scale. Is 'linear scale' too confusing? Or 'usual, linear scale'?

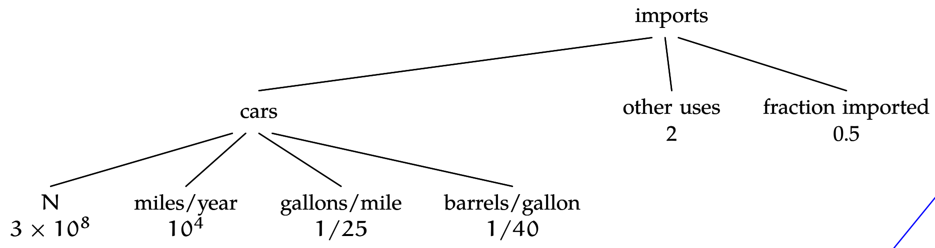

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3\times10^{8}\,\text{cars}\times\frac{10^{4}\,\text{miles}}{1\,\text{car–year}}\times\frac{1\,\text{gallon}}{25\,\text{miles}}\times\frac{1\,\text{barrel}}{40\,\text{gallons}}\sim3\times10^{9}\,\text{barrels/year}.$$

The two adjustment leaves contribute a factor of $2\times0.5=1$, so the import estimate is

$$3\times10^{9}\,\text{barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^{9} barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

Yes. They are always equal when the two numbers are the same. To prove this, set 2sqrt(nm)=n+m. Any n and m that satisfy this equation will have the two means be equal. The only n and m combinations that solve this have n=m

I found Problem 1.5 confusing. Both scales go up to infinity. So: what does "the midpoint on the log scale" denote? Are we talking about the midpoint of a line segment, measured on a log scale?

I think its talking about the idea that we have talked about in class for the number 10, for the geometric mean we use 3, but the midpoint on the usual scale is 5.5 (halfway between 1-10).

also, I'd switch 1.2 & 1.3 (keeping the same problem & using a side bar)...then you'd have the trees to help with the multiple estimation parts

i feel like this would've been an easier example to start with than CDs, because it requires numbers and intuition that more people are familiar with. the details of how CD data storage works seems a little esoteric to me...

yeah I agree, this example is a lot easier to follow and I have a better grasp of divide and conquer now.

I agree also, I like this example better than the CD. I feel like you need a certain amount of technical expertise to understand the CD example, but not with this one. I think this example could be more beneficial to more people.

i also agree, capacity seems a lot more vague than oil and this has less methods of calculation so it's a simpler example to start with (whereas the cd used many different approximations)

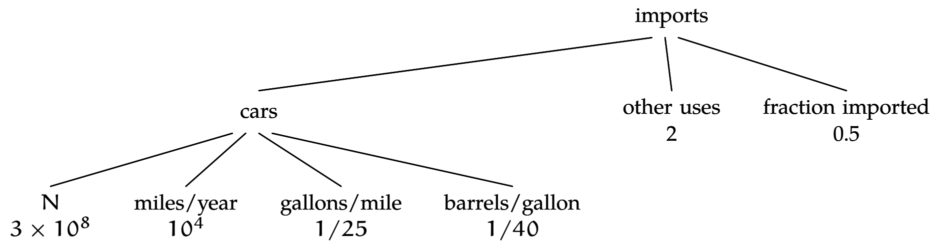

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8\,\text{cars} \times \frac{10^4\,\text{miles}}{1\,\text{car–year}} \times \frac{1\,\text{gallon}}{25\,\text{miles}} \times \frac{1\,\text{barrel}}{40\,\text{gallons}} \sim 3 \times 10^9\,\text{barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

3×10^9 barrels/year.

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

> **Problem 1.5 Midpoints**
> The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

Global comments

I'm having a hard time reading this section (I have no background really in coding, nor do I use UNIX). I feel like I'm going to have to read this a bunch of times before I begin to understand what is actually going on. I understand the use of divide in conquer in this context, I'm just mostly confused following these examples (considering how easy and great the previous sections were to read).

I agree with this. There is so much technical jargon in this section that most of my time is spent trying to understand and remember the new terms/commands instead of divide and conquer.

Agreed - I suggested there be a short description of how UNIX works preceding this chapter. I don't have much UNIX experience either and was quite confused, so I imagine everyone else in the same boat would have similar problems.

I feel the exact same way....divide and conquer permeates all of this section, but that's the extent of my understanding.

I agree with previous comments; it would also be helpful if you defined or explained some of the terms being used in this section – utility, input/output, like you did for "filter"

was UNIX the first program with this kind of search function (ie. searching for trivial, nontrivial, trivializes along with trivial)? This seems to be pretty ubiquitous with all searches nowadays.

> Wikipedia says that there was earlier precedents. But UNIX seems to be the first system to use regular expressions everywhere. A nice feature of the UNIX approach is that the grep program gives you back lines, on which you can do more things (like in class when we extracted the email addresses from the Stellar HTML file).

This section requires a foundational knowledge of UNIX that I don't think most people have...the overall message of divide and conquer is being lost amongst all these functions that are being used.

I agree to a certain extent. While I'm familiar with Unix, I could imagine this section discouraging those who do not. Is it possible to illustrate the concepts another way?

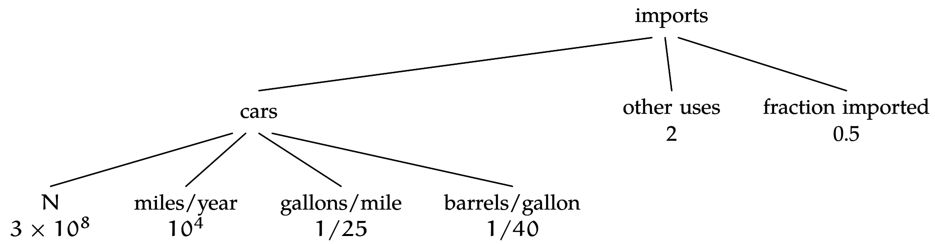

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3 \times 10^9 \text{ barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

> **Problem 1.5 Midpoints**
> The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

I think the UNIX example not broad enough for someone unfamiliar with the UNIX system to understand quickly. I like examples that are much more general/things everyone encounters on a fairly regular basis much better.

> I agree with this comment completely. I could follow the example because I have some experience with UNIX. However, I feel like this would have been fairly hard to follow for someone with no UNIX experience, especially since a large part of the people in this class are not course 6.

I think the unix example is a good example of how divide and conquer can be used for more complex problems and how it is applicable outside the subject matter. However, I think that the heavy programming language necessary to make this point distracts readers from the main messages. This chapter seems more like an in depth example than a lesson. Is that what you were going for?

Is that why the tree diagrams focus on one thing in the larger spectrum of example? For example, when calculating the imported oil consumption in the U.S., the estimation was also based on the barrals of oil used for cars.

I have no idea what this means....hold the user captive in their pre-designed set of operations? How is that related to email or writing a document?

Ok. I think this has finally made sense. I took Linguists and it definitely makes it easier to model and understand this tree diagrams. As well as the ones we make via actual calculations using divide and conquer

Unlike the other divide and conquer problems seen in this class, this method applied to Unix in this specific example does not seem very useful. The question of having a backwards dictionary seems improbable in real life.

Will we be using other methods to solve this Unix problem over the course of the class?

I don't really like this as a largescale example, I feel like I need to know more about UNIX to understand it better.

I had a difficult time understanding this article. I feel that the last example was much more straightforward. This seems like it requires a decent level of knowledge of UNIX to understand anything more than the very basic idea that large problems need to be broken into smaller problems.

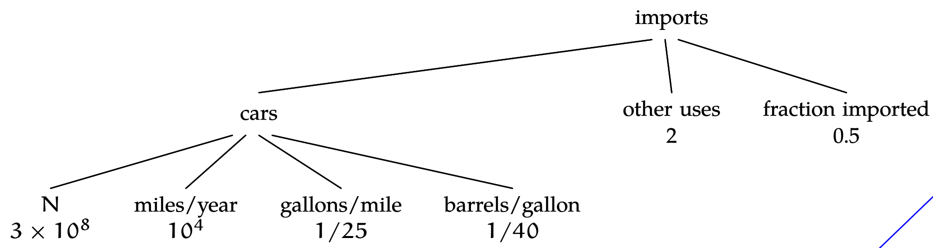

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3 \times 10^9 \text{ barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

Read Section 1.5 for Tuesday's lecture (due to the Monday holiday). The memo is due by Monday at 10pm. Don't worry about the last two pages – they are a bunch of end-of-chapter problems (but some of them appear on your first homework).

This example really hit home what trees and divide and conquering were, how to use them, and the thinking process of breaking down big problems into manageable pieces. Because this example did not include "speaking to your gut" it was much more digestable and I would have liked to have seen this example proceed the estimation divide and conquering problems.

I also feel this way. I supposed partly because this is a course 6 class, and we are mostly engineers, that tis kind of example would really hit home with us.

I like your use of "speaking to your gut" and "digestible"

What makes one operating system "a cousin" of another? Are they simply operating systems that were created by the same people? Or those with similar features? Is there any formal classification?

Does an example not actually about estimation really have place in a book about estimation? I believe you made your point about how useful divide and conquer is when you first introduced the idea.

Is it truly portable? (I don't know anything about computers.) For example, with regards to web browsers, I've heard that either Chrome of Firefox is portable, but some of its features take data stored in IE, since it is assumed that all computers have IE. Are there conditions/pre-existing features necessary for UNIX/GNU/Linux to run?

It's important to keep in mind that "most successful" is still not particularly successful in comparison to corporate owned operating systems, due to both marketing and ease of use.

The point here is just that Linux is the most successful variant of UNIX, and it's not due to marketing. Your comment depends on your definition of success. I wouldn't want most research servers running Windows, even if it's a useful OS for home and business users.

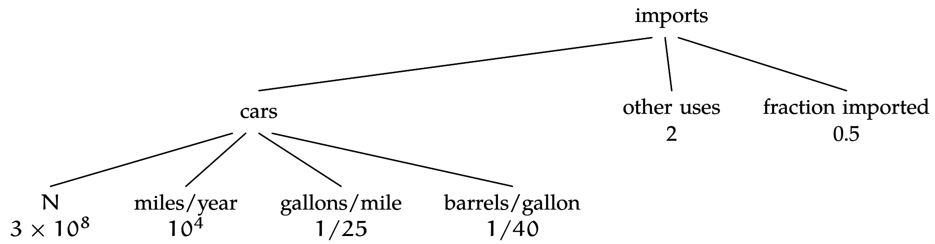

All the leaves have values, so I can propagate upward to the root. The main operation is multiplication. For the 'cars' node:

$$3 \times 10^8 \text{ cars} \times \frac{10^4 \text{ miles}}{1 \text{ car–year}} \times \frac{1 \text{ gallon}}{25 \text{ miles}} \times \frac{1 \text{ barrel}}{40 \text{ gallons}} \sim 3 \times 10^9 \text{ barrels/year}.$$

The two adjustment leaves contribute a factor of $2 \times 0.5 = 1$, so the import estimate is

$$3 \times 10^9 \text{ barrels/year}.$$

For 2006, the true value (from the US Dept of Energy) is 3.7×10^9 barrels/year – only 25 higher than the estimate!

Problem 1.5 Midpoints
The midpoint on the log scale is also known as the geometric mean. Show that it is never greater than the midpoint on the usual scale (which is also known as the arithmetic mean). Can the two midpoints ever be equal?

1.5 Example 4: The UNIX philosophy

The preceding examples illustrate how divide and conquer enables accurate estimates. An example remote from estimation – the design principles of the UNIX operating system – illustrates the generality of this tool.

UNIX and its close cousins such as GNU/Linux operate devices as small as cellular telephones and as large as supercomputers cooled by liquid nitrogen. They constitute the world's most portable operating system. Its success derives not from marketing – the most successful variant, GNU/Linux, is free software and owned by no corporation – but rather from outstanding design principles.

These principles are the subject of *The UNIX Philosophy* [13], a valuable book for anyone interested in how to design large systems. The author

Is this also true of UNIX, or is UNIX actually owned by a corporation?

As I understand it, Unix, while an owned trademark, refers mainly to a standardization. Whereas Linux is an actual open-source OS.

I think UNIX was developed partially at Bell Labs

I believe Unix IP has ownership; that is in fact the impetus behind Linux development. An apocryphal(?) story is that it was named because it was a "Unix Like" kernel.

I believe the Apple OS is also UNIX-based (though Windows is not)

I'm also a little unsure of how UNIX works and could use a little more info about how it is related to GNU/Linux as cousins.

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

In a discussion of estimation, why include the philosophy points that don't pertain to divide & conquer? You could address the 4 related points first and then mention the others, but having some with comments and some without seemed a bit jarring to me.

I also found them somewhat distracting, especially because without some explanation, I really have no idea what they mean. Each of them raises questions instead of furthering my understanding of divide and conquer.

I agree, the inclusion of the other 5 points is unnecessary and confusing. I found myself constantly evaluating whether I was currently reading one of the 4 important points. Also, to those not from a computer science background, these other points would only confuse them, even though they're not important to the later discussion.

I think the other 5 tenets were included for completion (why leave out parts of a set?), but I agree that for clarity they should be dropped.

Is this itself an example of dividing and conquering?

I would say no in the estimation since, as there is no number problem being solved with these techniques; however, as has been stressed many times, these skills aren't just useful for the purposes of guestimating numbers. It can be applied to a variety of approaches to thinking, in which case, yes, it is?

No, it's identifying general themes about UNIX, so it's more about lumping common aspects together. The 9 concepts aren't separate parts that you could put together to make UNIX.

Philosophy related to divide and conquer? That seems a bit of an odd parallel. Maybe structure makes more sense than philosophy?

I don't think this is limited to d&c. It's not a structure, it's a set of design principles, which can be referred to as a (design) philosophy.

What do you mean by comments? It seems like all of them have comments.

I think 3 through 7 are just the tenets, without comments. I agree that it's not particularly clear what constitutes the added comment. I think though, that just the first sentence of each item is the tenet itself.

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

The lack of memory in early UNIX systems also played a large factor in the division of programs into smaller pieces as each could only access something on the level of 64K of memory - not enough to do complex operations as a single program but requiring multiple programs with very specialized function. From other UNIX history readings I've had for other classes it seems less of a "UNIX made for divide and conquer" and more of a "divide and conquer happens to work well with UNIX"

I would like to see the actual principles set apart from the commentary somehow. Perhaps making the principles bold would do the trick?

agreed.

And/or having a line break after the principle so it's on its own line before the comments.

I think this is a very useful principle in divide and conquer that should have been stated at the very beginning. It helps me to think about what sort of categories to divide into to make the estimation easier for me.

I understand that this is the way to go. The hard part is realizing you know something about things you don't have direct knowledge about. I found this the most frustrating part of the diagnostic.

I agree, this could be the first sentence in the divide and conquer section. I have talked to people in the class and thought that you knew a lot of facts they wouldn't when the important part is that you need to isolate things you know.

Should this actually read Section 1.3?

Yeah, typo most likely.

This is information that could have been stated earlier when describing "divide-and-conquer"

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

Would this also improve efficiency, because now the computer can search in parallel instead of in serial for each leave? Or no?

Hmm I'm curious about this as well - would it make things faster too?

Computers can only search in parallel, if they have more than one designated processor. These days, you hear about "dual-core" or "quad-core" processors. Those can do parallel computing. Otherwise, smaller programs won't make computing faster. They'll just make it easier for you to debug and write the overall program.

In fact, having more, smaller programs could slow things down as more things would have to be loaded into memory, taking longer than one longer single load.

This is similar to constraint based design in Mechanical engineering. The idea is you choose the degrees of freedom you don't want and then only place constraints that kill the individual degrees of freedom. This way you're not overconstrained which often leads to mechanical failure.

This is the idea of modularity, right? If we keep things simple and contained it will be easier to work with.

I think it would be nice to state modularity in the description

Modularity doesn't just work on this simplest level, but also at higher levels where parts are connected. Connections and independence between parts are also key to modularity. You can have unitaskers that aren't interchangeable. I agree that it might be a theme, but it's not this one exactly.

use of modularity in programming

is this saying that we should make clear trees? I am not getting the connection to divide and conquer

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

I'm not sure I understand what the difference between 1. and 2. is? Both seem to break up a large task into small tasks.

I think one is about partitioning where things go and how you find them (1) while the other is about actually running scripts to do tasks, like work processing. At least, thats what I understood.

The way I saw it was: 1. Make the programs small. 2. Make the programs simple. I guess they could be combined, but for some reason my intuition tells me that it's easier to understand as two separate points.

Yeah, I agree. The way I see it, in the first point the author is saying break complex tasks into smaller ones, so each individual task is easier to understand and solve. And then in the second point he's saying specifically to make each program do only one task, and that this way, the program is easier to debug

In addition to being simpler and easier to debug, the main point is that since its one task, you can clear you mind and make that one task done well. A small program is fine, but if it still isn't done well then its pointless. After making a bunch of small program that work very effectively, chances are when you put them together they will work effectively as well.

Another thing that is nice about small, simpler specialized programs is that they can be used in many contexts. With the example of spell-checking, it can be used by more than one application (Word, email etc)

How does this apply to the need for more processes to be run concurrently?

Do you mean more processes need to be run to accomplish a complex task? That could be a problem, as stated in an earlier comment, when it comes to processing power, but the small processes wouldn't be running concurrently, they'd be running in series, which may cause a slowdown.

I didn't understand which of the nine points were the four incorporating divide and conquer until I read the comments. Even then, I had to assume that the four with explanations were divide and conquer examples

Go into more detail- how is this using divide and conquer?

I don't know if this is supposed to be one of the four, but I think this is roughly applicable if you substitute "prototype" with "constructing problem solving formulas"

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

Prototypes are usually for testing a particular piece of your idea, independent of the rest of the structure, which is the fundamental idea behind divide and conquer. This point should be elaborated a bit more.

I don't really understand how this is part of a philosophy. isn't it the equivalent of don't procrastinate?

No, I think that what they mean is that it's more important to have a prototype out that's 70% done in 2 weeks than an almost finalized product that's 95% done in 2 months because user response is valuable in redirecting projects and allows for a more intimate iterative process. It's less about procrastination and more about understanding that it's important to get a prototype out there asap for the users to say what they like and dislike.

I took this to be a philosophy for the developer to get off the ground, rather than getting something out there to be tested by users. I interpreted it as how it's much easier for me to edit something I wrote poorly than write something new. Therefore, I should just write something rather than stare at the blank paper.

in upop, we learned something very similar to this concept, and called it, "make mistakes faster"

Hopefully an iterative spiral model, with prototypes being built starting from lower fidelity prototypes to higher fidelity.

This is also applicable to rough estimates: ease over precision

Sort of, but aren't we all about efficiency (rounding to 1, few, and 10) AND portability (perhaps using pi*10ˆ7 to remember seconds/yr)?

yes but if we had to choose between the two, we would choose portability; because that is the whole premise of estimation.

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

I'm not sure I understand this, in the context or UNIX or in the context of divide and conquer.

I also am not so sure what they mean about choosing portability over efficiency. Could someone please explain this?

does portability mean less code rather than more efficient code?

I'm not really a coder, but my guess was that it's saying that it's better to have very clear, defined code, even if that means adding a couple extra lines.

yeah the use of portability is confusing.

I think this also goes along with being able to use the same code with multiple programs (back to spell checking, using it for word processing, email etc) instead of having each program have their own spell-checker (or whatever it is they might share)

I thought this concept of being able to use the same code with multiple programs was called modularity? But I also don't have a strong computer science background and might be wrong. But the post regarding extra lines for clear definition makes sense too.

I think portability is being mentioned for its benefit of being less specific to certain platforms, etc. Choosing simpler functions or methods more universal is more accepted than exploiting the most efficient approach on a given platform.

I also think that portability may be referring to modularity. Being able to you a small code for multiple things. Although, im not sure how they justify it being better than efficiency.

A tool that you can use for many different problems, that will get you close to the answer is much better then one that will give you one exact answer to only one problem.

I think this refers to the speed which is required in solving approximation problems. You want to come up with a plan of attack as soon as possible.

Who wrote these philosophies? Did a bunch of programmers came together and write down what they thought was important at the time?

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

What is a "flat" text file? Does this imply using .rtf instead of something more complicated, or storing things just once instead of layering them over older versions of the same file?

I think it means plain ASCII files. Even RFT files have some formatting in them, and so require programs with special abilities to read. But any computer ever built can open and display an ASCII file readably without any special software.

I agree with the second post - I think a "flat" text file is basically one that doesn't require specialized software to open (RTF, etc).

i wouldn't have known this either, but it seems to not matter in this context of divide-and-conquer.

Just for the sake of discussion, I think "flat" can also be used to distinguish a file that can be read and understood sequentially (often ASCII) from a database system that has a non-standard underlying file/storage structure.

Even though these aren't relevant to divide and conquer I still might like to see a sentence describing what is meant by each.

I feel that the reasons it is not very easy to understand this it's because very little has to do with Divide and Conquer. reading the next pages helps to understand it better

What is Flat?

more details about these would be helpful

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

What is software leverage?

I also have no idea. Maybe simply that there is a lot of free software out there that one can incorporate into UNIX?

I was unclear about the meaning also. After looking it up, I found a pretty good explanation of it here: http://books.google.com/books?id=MFzVYQsD60AC&lpg=PA67&o
xiEl9bI&dq=%22software%20leverage%22%20definition&pg=PA67#v=onepage&c
. It is a bit of a read and would have been better had the term been explained here.

I didn't understand a lot of the phrases in 1-9. Perhaps an easier concept on something less technical, or more definitions. I understand the divide & conquer parts, but thinking about the new words made it difficult to not get overwhelmed.

I agree, I understand the divide and conquer parts, but I had to look up a lot of these terms. Maybe another sentence in 1-9 that explains these terms of gives an example would be helpful

Or just omitting 3 through 9? They don't seem to contribute much to the topic at hand, and many of us seem to be confused by them.

I'm not familiar with this terminology. What is "software leverage"? And how do you use it?

Agreed - this is confusing for me as well.

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

I'm not sure what a shell script is - since it seems that there's some confusion over some UNIX specific points, perhaps it would be best to open this section with a short description of how UNIX works?

Agreed; this section so far isn't so accessible if you don't know much about operating systems or programming, so maybe a bit more explanation would be nice (although maybe then it'd defeat the purpose of using this as a parallel to explain divide and conquer and just turn into a lesson on UNIX which isn't what you want either...)

Agreed. I looked up shell scripts on wikipedia: A shell script is a script written for the shell, or command line interpreter, of an operating system. It is often considered a simple domain-specific programming language. Typical operations performed by shell scripts include file manipulation, program execution, and printing text."

Shell scripts may be useful because they are typically portable within UNIX based operating systems. Without incorporating dependencies to specific software (ex. using flat txt files), the code is more likely to work on other UNIX based operating systems.

I thought a captive UI was sort of UI that disallows interactivity once the user starts the application...but I'm not sure how email is like that...

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

What is a "captive user interface"? (it just doesn't sound pleasant...) Because of that, I don't really understand what is going on in this tenet.

Once again...agreed, I have no idea what this means either.

> p.s. If you use the "I agree" option, then this question (which clearly many people have) will get moderated upward and will be one of the first that I see when I look at all responses.
>
> Whereas NB is not (yet?) smart enough to do the same if you verbally agree. AI has a ways to go. In short, use the "I agree" option (or "I don't agree") if it captures what you want to say.

I'm not entirely sure but based on the rest of the paragraph I would guess they are limited GUI interfaces that don't allow maximum user freedom. Programs like MATLAB and photoshop (ignoring common filters) are what I would consider good examples of a non captive user interface, because they allow the user to define and modify operations to their liking.

I think "captive user interfaces" refer to super fancy layouts but simply for the glittery oo ahhh effect. Actually come to think about it, its probably referring to restrictive UIs, like the old DOS programs.

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

You're right, it's not pleasant. Though it's common. For example, the typical word processing program allows only one way to edit and create your document: You must use its user interface or else. For example, how else can you make an OpenOffice or MS-Word document, except by using the corresponding program?

Whereas with a non-captive-UI, you could use any method you like to make a word-processing file, which is then (say) turned into PDF by the word processor. TeX is an example of this kind of program. I have many programs that themselves write TeX code, which is then compiled to make PDF.

The trees in this chapter are made by a similar method. The figure-drawing program is called MetaPost, which is much easier to program than directly writing PostScript. But even so, making the divide-and-conquer trees become a bit tedious because I kept repeating so much of the MetaPost code. So I wrote a that program takes a simple text-file representation of a tree and converts (compiles) it into MetaPost, which is then compiled into PostScript and PDF. That would be very hard to do using a graphical figure-drawing program. Because of the captive user interface, only a human typing and clicking at the keyboard and mouse can generate the drawing.

this clearly explains what captive user interface is. thanks for the examples. I think the ones you listed here are a little better than the ones in the reading.

Are any examples of this? I can understand that user interfaces are large and hard to debug, but how exactly would you divide and conquer this?

Again, some unfamiliar terminology.

I don't know what captive user interfaces are.

It seems like this principle has been violated by quite a few programs nowadays...

I would be interested to see an example of such a program because I'm still not really sure that i understand what is meant here.

...yeah but shitty interface repels the average, non-programmer user

i don't really understand how these two concepts are related, or how the unix really solves the problems of the first paragraph.

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

I'm not sure if I'm understanding this correctly, but doesn't leaving the program from the user to define tend to deter the user from using the program since it means extra work?

> Not if you do it right, I think its good to give the user some control over the processes. As long as the smaller tasks are atomic the user should be able to do what they want without messing things up.

So they don't use UI? Or they have separate UIs for each program that can be combined into one?

> I think it might be more that they are not limited to the buttons seen in like, typical email. They can either make their own or use others that the original designers didn't think of.
>
> > I agree... its not necessarily that a UI doesn't exist, but rather that if a UI module does exist, it doesn't limit functionality. Simple programs exist first, and a UI may be built on top of it to provide visual support, but may or may not incorporate all of the program's functionality / ability.

I'm not sure I completely follow this. I'm not familiar with the terms, but it seems like the author is being a bit repetitive. Points 1 and 2 were pretty similar–saying to split programs into smaller programs. Here, he's also saying the same thing–"solve complex tasks by dividing them into smaller tasks..."

This sounds to me more like robustness (not necessarily redundancy), and less like our goal of simplification when we use divide and conquer, right?

this is a much easier term to understand than the previous ones presented.

Works like a transfer function in systems and signals. If you can convert the elements into transfer functions, it's easy to design a system to match the desired performance requirements.

This is a superb definition of a program

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

I like how you defined filter here, but I feel like filter is actually a more-known term than the ones you did not define above.

I agree. I like the definition, but some others are more necessary.

I agree too. I needed definitions for almost every term on this list.

I agree, too. This filter example makes sense to me. Basically, make your programs all mesh with each other–input for one is the output for another, so you don't have to keep using user input all over

I don't quite understand how this definition of filter makes it a good leaf in a tree...

Don't take "leaf" so literally; you may be thinking of leaves as the end of the branches, but I think most often leaves just mean another branching point. If you think of it that way, filters take data, process it, and then out put data for input into the next branch just like all the other branches we have talked about in other trees in class.

No – a leaf necessarily (by definition) has no children. I think the right term here would be "internal node", not leaf.

I agree – this doesn't make much sense the way it is written. The idea we seem to be going for is that, for example, multiplication and dimensional analysis combine easily in arbitrary ways, just like filters, but I'm not sure the way this is explained works.

Here you explain how filters relate to the idea of divide-and-conquer... but it's not so clear for some of the earlier points. I'm not sure if you're trying to just list the main tenets of the UNIX philosophy, or if each of these points is supposed to correlate to some aspect of divide and conquer.

It says in the line proceeding the list that only the 4 with explanations are relevant. Maybe including the others is unnecessary since they don't seem to have any real relevance to divide an conquer.

Agreed. Maybe some explanation of UNIX followed by just the tenets that relate to d&c would be clearer.

I think this is a very good informative section; I personally really enjoy sections like this because it's the little facts that make one a more worldly, learned individual. That said, it does seem to be a lot of info that isn't immediately relevant to estimation? On the other hand, I prefer it this way and I'm sure most MIT (and college) students do to, and that's the target audience.

isolates nine tenets of the UNIX philosophy, of which four – those with comments in the following list – incorporate or enable divide-and-conquer reasoning:

1. Small is beautiful. In estimation problems, divide and conquer works by replacing quantities about which one knows little with quantities about which one knows more (Section 7.3). Similarly, hard computational problems – for example, building a searchable database of all emails or web pages – can often be solved by breaking them into small, well-understood tasks. Small programs, being easy to understand and use, therefore make good leaf nodes in a divide-and-conquer tree (Section 1.3).
2. Make each program do one thing well. A program doing one task – only spell-checking rather than all of word processing – is easier to understand, to debug, and to use. One-task programs therefore make good leaf nodes in a divide-and-conquer trees.
3. Build a prototype as soon as possible.
4. Choose portability over efficiency.
5. Store data in flat text files.
6. Use software leverage to your advantage.
7. Use shell scripts to increase leverage and portability.
8. Avoid captive user interfaces. Such interfaces are typical in programs for solving complex tasks, for example managing email or writing documents. These monolithic solutions, besides being large and hard to debug, hold the user captive in their pre-designed set of operations.

   In contrast, UNIX programmers typically solve complex tasks by dividing them into smaller tasks and conquering those tasks with simple programs. The user can adapt and remix these simple programs to solve problems unanticipated by the programmer.
9. Make every program a filter. A filter, in programming parlance, takes input data, processes it, and produces new data. A filter combines easily with another filter, with the output from one filter becoming the input for the next filter. Filters therefore make good leaves in a divide-and-conquer tree.

how is a program different from a filter otherwise? don't filters remove noise? but that's not the definition used here.

A filter seems to be a type of program, that may or may not remove noise. A filter, as it is used here, seems to be more of an operation.

I think you are thinking of a filter too literally, not just signal processing, but for example, a text editor would have a small program(more likely a function) that takes input from the keyboard and displays it on the screen. It filters the keystrokes into ASCII characters/pixels on your screen. I think that's what's meant by filter...

I agree. It would be very useful if you identified the four principals that apply to both programming and divide and conquer. Also, those four points should be parallel to one another in their explanations.

This section is hard to understand for someone with little to no programming knowledge. I'd like if more terms were defined or if programming was compared to divide and conquer as opposed to the opposite (which is how it seems now)

As examples of these principles, here are two UNIX programs, each a small filter doing one task well:

- `head`: prints the first lines of the input. For example, `head` invoked as `head -15` prints the first 15 lines.
- `tail`: prints the last lines of the input. For example, `tail` invoked as `tail -15` prints the last 15 lines.

How can you use these building blocks to print the 23rd line of a file?

This problem subdivides into two parts: (1) print the first 23 lines, then (2) print the last line of those first 23 lines. The first subproblem is solved with the filter `head -23`. The second subproblem is solved with the filter `tail -1`.

The remaining problem is how to hand the second filter the output of the first filter – in other words how to combine the leaves of the tree. In estimation problems, we usually multiply the leaf values, so the combinator is usually the multiplication operator. In UNIX, the combinator is the pipe. Just as a plumber's pipe connects the output of one object, such as a sink, to the input of another object (often a larger pipe system), a UNIX pipe connects the output of one program to the input of another program.

The pipe syntax is the vertical bar. Therefore, the following pipeline prints the 23rd line from its input:

```
head -23 | tail -1
```

But where does the system get the input? There are several ways to tell it where to look:

1. Use the pipeline unchanged. Then `head` reads its input from the keyboard. A UNIX convention – not a requirement, but a habit followed by most programs – is that, unless an input file is specified, programs read from the so-called standard input stream, usually the keyboard. The pipeline

   ```
   head -23 | tail -1
   ```

Comments on page 3

I'm a bit confused here, and in the following section: are we writing this somewhere, like code, or is this an already made program like, say, notepad?

head and tail are already programs that exist in UNIX, like in Athena the ls and cd commands that you are probably are familiar with. When they are combined below that is written by the user.

Would run be a simpler term?

This example for filters is helpful, but I feel like the author's point (#9) about filters was already the clearest. An example about #8 (captive user interfaces) would have been helpful too

I disagree I think the verbal explanation of a filter was pretty clear but this makes sure I understand it 100% and I'm glad it's here. I do agree that examples of the other ones would be very helpful

This looks a little strange. Maybe put in quotations the actual input, or make it a different font, or something.

How you arranged it below, by putting it on a separate line, would work well too

I would either put it in italics or make it a different color

I think it already is in a different font from what I can tell, but a different color would work wonders.

he does use a different fond...however, it would be useful to make a side note for people who don't code about standard practices...courier is always used to designate what you would actually type

So far this whole example is disorganized and a bit hard to follow. Overall it needs to be a bit better explained. You should probably just choose a few points and elaborate on them, rather than just listing everything you can think of.

I really do not understand what this is talking about at all. it prints the first 15 lines of what

As examples of these principles, here are two UNIX programs, each a small filter doing one task well:

- `head`: prints the first lines of the input. For example, `head` invoked as `head -15` prints the first 15 lines.
- `tail`: prints the last lines of the input. For example, `tail` invoked as `tail -15` prints the last 15 lines.

How can you use these building blocks to print the 23rd line of a file?

This problem subdivides into two parts: (1) print the first 23 lines, then (2) print the last line of those first 23 lines. The first subproblem is solved with the filter `head -23`. The second subproblem is solved with the filter `tail -1`.

The remaining problem is how to hand the second filter the output of the first filter – in other words how to combine the leaves of the tree. In estimation problems, we usually multiply the leaf values, so the combinator is usually the multiplication operator. In UNIX, the combinator is the pipe. Just as a plumber's pipe connects the output of one object, such as a sink, to the input of another object (often a larger pipe system), a UNIX pipe connects the output of one program to the input of another program.

The pipe syntax is the vertical bar. Therefore, the following pipeline prints the 23rd line from its input:

```
head -23 | tail -1
```

But where does the system get the input? There are several ways to tell it where to look:

1. Use the pipeline unchanged. Then `head` reads its input from the keyboard. A UNIX convention – not a requirement, but a habit followed by most programs – is that, unless an input file is specified, programs read from the so-called standard input stream, usually the keyboard. The pipeline

   ```
   head -23 | tail -1
   ```

The notation of this example confuses me.

The 'head -15' and 'tail -15' refer to text that would be typed into a command prompt or terminal.

I think the argument syntax for this function may confuse people not familiar with command line or terminal commands. I think a sentence or two describing the syntax, or at least making it less confusing to others might help.

I think these small examples would be better suited within the points that precede it, or at least identify the principals they match up with.

From solving this problem, I would guess that search engines use this divide and conquer method in their algorithms.

This is a really good problem, we did similar problems in 6.033 at around the same time

A page and a half was spent describing the logisitics of UNIX commands and how to use them to answer this question. This section is supposed to be describing how divide and conquer was used when UNIX was designed, and instead my time is wasted with a description of how to write UNIX commands, which I already know how to do. I understand that not everyone who reads this knows UNIX and UNIX commands, but in that case, I don't think this is a good example if so much explanation is necessary.

I liked this example and I think the time spent to explain the pipe combinator was worth it. Now I can think of the lines connecting nodes of a tree as the pipe operator and the node itself as the filter.

I think instead of listing these long commands that non course 6exers dont know much about, you could do more examples of the "print" one. There are many many problems using lists that could be relevant to divide and conquer

what do you mean? print only 23 lines? or print the 23rd line in a larger series?

This is a really good example, easy to understand for people who haven't had programming experience before.

As examples of these principles, here are two UNIX programs, each a small filter doing one task well:

- `head`: prints the first lines of the input. For example, `head` invoked as `head -15` prints the first 15 lines.
- `tail`: prints the last lines of the input. For example, `tail` invoked as `tail -15` prints the last 15 lines.

▶ *How can you use these building blocks to print the 23rd line of a file?*

This problem subdivides into two parts: (1) print the first 23 lines, then (2) print the last line of those first 23 lines. The first subproblem is solved with the filter `head -23`. The second subproblem is solved with the filter `tail -1`.

The remaining problem is how to hand the second filter the output of the first filter – in other words how to combine the leaves of the tree. In estimation problems, we usually multiply the leaf values, so the combinator is usually the multiplication operator. In UNIX, the combinator is the pipe. Just as a plumber's pipe connects the output of one object, such as a sink, to the input of another object (often a larger pipe system), a UNIX pipe connects the output of one program to the input of another program.

The pipe syntax is the vertical bar. Therefore, the following pipeline prints the 23rd line from its input:

```
head -23 | tail -1
```

But where does the system get the input? There are several ways to tell it where to look:

1. Use the pipeline unchanged. Then `head` reads its input from the keyboard. A UNIX convention – not a requirement, but a habit followed by most programs – is that, unless an input file is specified, programs read from the so-called standard input stream, usually the keyboard. The pipeline

```
head -23 | tail -1
```

I am confused about this, wouldn't the tail -1 be printing the 24th line?

You should think of the line as saying "tail -1 of head -23" which would mean "print the last line of the 1st 23 lines. "

Agreed, it seems as though we're printing the first 23 lines then the final line of the entire file.

the idea is that 'tail -1' is run on the 'head -23' output. There isn't a 24th (or greater) line at that point.

He talks about feeding outputs to other functions later. That should clear it up.

We don't always multiply, right?

A lot of the time it's just dimensional analysis, so it is pretty much multiplying by different unities.

I don't like the way this is introduced. I think it would be better to suggest it is a metaphor. Something like: "you can think as the filters as a pipe."

I don't like the way this is introduced. I think it would be better to suggest it is a metaphor. Something like: "you can think as the filters as a pipe."

I don't like the way this is introduced. I think it would be better to suggest it is a metaphor. Something like "you can think as the filters as a pipe."

this seems very intuitive to me. it was already explained above

I don't like the way this is introduced. I think it would be better to suggest it is a metaphor. Something like: "you can think as the filters as a pipe."

I don't like the way this is introduced. I think it would be better to suggest it is a metaphor. Something like "you can think as the filters as a pipe."

I don't like the way this is introduced. I think it would be better to suggest it is a metaphor. Something like "you can think as the filters as a pipe."

I thin

As examples of these principles, here are two UNIX programs, each a small filter doing one task well:

- `head`: prints the first lines of the input. For example, `head` invoked as `head -15` prints the first 15 lines.
- `tail`: prints the last lines of the input. For example, `tail` invoked as `tail -15` prints the last 15 lines.

▶ *How can you use these building blocks to print the 23rd line of a file?*

This problem subdivides into two parts: (1) print the first 23 lines, then (2) print the last line of those first 23 lines. The first subproblem is solved with the filter `head -23`. The second subproblem is solved with the filter `tail -1`.

The remaining problem is how to hand the second filter the output of the first filter – in other words how to combine the leaves of the tree. In estimation problems, we usually multiply the leaf values, so the combinator is usually the multiplication operator. In UNIX, the combinator is the pipe. Just as a plumber's pipe connects the output of one object, such as a sink, to the input of another object (often a larger pipe system), a UNIX pipe connects the output of one program to the input of another program.

The pipe syntax is the vertical bar. Therefore, the following pipeline prints the 23rd line from its input:

```
head -23 | tail -1
```

But where does the system get the input? There are several ways to tell it where to look:

1. Use the pipeline unchanged. Then `head` reads its input from the keyboard. A UNIX convention – not a requirement, but a habit followed by most programs – is that, unless an input file is specified, programs read from the so-called standard input stream, usually the keyboard. The pipeline

   ```
   head -23 | tail -1
   ```

This section focuses a lot on programming and not much on divide and conquer. I'd like to see divide and conquer more incorporated in this section, otherwise I feel like I'm getting lost in programming and missing the whole point of this section.

I like the analogy

I feel like I'm missing what this has to do with estimation. It makes sense and is interesting and all of those good things, but I'm sort of wondering how it's relevant.

It's great that this part is so clear and easy to follow for people who don't understand UNIX, and I think it points out the slight inaccessibility of the previous section, when there were a lot of terms the UNIX-unsavvy couldn't follow.

As a before noted UNIX-unsavvy person, I agree. This is getting much clearer, although going over the principles in class would also help a lot. Please don't forget this is also registered as a course 2 class.

I agree. Sometimes I have trouble reading this book because I get distracted with unfamiliar concepts. But when you include definitions, the readings become clearer and I see how it relates to the problem at hand.

I agree with the comments from the others about this being well explained. I have no experience with UNIX at all, but this is nice and clear.

It's funny - in 6.033 we're reading the actual UNIX paper, and your writing is making it 10X more accessible.

I see how this models the tree and divide and conquer, but I don't see what we are going to be estimating

I thought that these printed in quantities of 15? or if not what was the 15 for

just curious, can you do head -24 | tail -2 and get the same result, or does the tail only print the LAST line of the input

I believe this would print line 23 AND line 24, and we were aiming to print only line 23.

As the previous reply says, you would get line 23 and 24 since you are saying the last TWO lines of the input.

As examples of these principles, here are two UNIX programs, each a small filter doing one task well:

- `head`: prints the first lines of the input. For example, `head` invoked as `head -15` prints the first 15 lines.
- `tail`: prints the last lines of the input. For example, `tail` invoked as `tail -15` prints the last 15 lines.

How can you use these building blocks to print the 23rd line of a file?

This problem subdivides into two parts: (1) print the first 23 lines, then (2) print the last line of those first 23 lines. The first subproblem is solved with the filter `head -23`. The second subproblem is solved with the filter `tail -1`.

The remaining problem is how to hand the second filter the output of the first filter – in other words how to combine the leaves of the tree. In estimation problems, we usually multiply the leaf values, so the combinator is usually the multiplication operator. In UNIX, the combinator is the pipe. Just as a plumber's pipe connects the output of one object, such as a sink, to the input of another object (often a larger pipe system), a UNIX pipe connects the output of one program to the input of another program.

The pipe syntax is the vertical bar. Therefore, the following pipeline prints the 23rd line from its input:

```
head -23 | tail -1
```

But where does the system get the input? There are several ways to tell it where to look:

1. Use the pipeline unchanged. Then `head` reads its input from the keyboard. A UNIX convention – not a requirement, but a habit followed by most programs – is that, unless an input file is specified, programs read from the so-called standard input stream, usually the keyboard. The pipeline

   ```
   head -23 | tail -1
   ```

For me this is where the divide and conquer analogy gets a little lost, although after thinking about it I suppose you could say defining the input is just another level of the tree. You might want to make that connection clearer.

good to know

So does this notation just mean that the output of head -23 is used as the input of tail -1?

Yes, thus the pipeline analogy.

So this will print only the 23rd line right?

therefore reads lines typed at the keyboard, prints the 23rd line, and exits (even if the user is still typing).

2. Tell head to read its input from a file – for example from an English dictionary. On my GNU/Linux computer, the English dictionary is the file /usr/share/dict/words. It contains one word per line, so the following pipeline prints the 23rd word from the dictionary:

```
head -23 /usr/share/dict/words | tail -1
```

3. Let head read from its standard input, but connect the standard input to a file:

```
head -23 < /usr/share/dict/words | tail -1
```

The < operator tells the UNIX command interpreter to connect the file /usr/share/dict/words to the input of head. The system tricks head into thinking its reading from the keyboard, but the input comes from the file – without requiring any change in the program!

4. Use the cat program to achieve the same effect as the preceding method. The cat program copies its input file(s) to the output. This extended pipeline therefore has the same effect as the preceding method:

```
cat /usr/share/dict/words | head -23 | tail -1
```

This longer pipeline is slightly less efficient than using the redirection operator. The pipeline requires an extra program (cat) copying its input to its output, whereas the redirection operator lets the lower level of the UNIX system achieve the same effect (replumbing the input) without the gratuitous copy.

As practice, let's use the UNIX approach to divide and conquer a search problem:

▶ *Imagine a dictionary of English alphabetized from right to left instead of the usual left to right. In other words, the dictionary begins with words that end in 'a'. In that dictionary, what word immediately follows* `trivia`*?*

This whimsical problem is drawn from a scavenger hunt [26]created by the computer scientist Donald Knuth, whose many accomplishments include the TEX typesetting system used to produce this book.

Comments on page 4

What does it mean that it reads lines from the keyboard? does this mean that after you type 23 lines of text it will execute the command?

I don't know much about programming and it's unclear what you mean by it reads lines typed at the keyboard. How does this differ from what we're looking for?

At the beginning of the problem we said we wanted the 23rd line of the file, so we probably want our filters to start with the contents of the file. Otherwise it would take whatever you are typing at the keyboard and return the 23rd line of that.

I'm not sure what the backslashes do

What is different between this and #2?

The difference between 2 and 3 clearly stated would be helpful.

I'm unclear on this. Is head going to read lines from the keyboard and then switch to the file? How does it know when to stop taking input from the keyboard?

What's the point of this if you can read from a file?

I'm not sure I fully understand this either. So does this put the input into a file, then read it from the file? So you're not directly reading the user input?

yeah i dont understand the nuances of this either

The wording is very technical and could use some more basic explanations about what it means to connect a file to an input.

While describing these three possible input methods for getting information into the pipe, I dont see what it adds to the divide and conquer section. It almost seems to be getting bogged down in the implementation of a specific case rather than the idea and effectiveness of divide and conquer.

I had the same comment while reading it. Although as a course-6-er I was interested in the explanation, it added nothing to the divide and conquer idea and obfuscated the reading.

therefore reads lines typed at the keyboard, prints the 23rd line, and exits (even if the user is still typing).

2. Tell `head` to read its input from a file – for example from an English dictionary. On my GNU/Linux computer, the English dictionary is the file `/usr/share/dict/words`. It contains one word per line, so the following pipeline prints the 23rd word from the dictionary:

```
head -23 /usr/share/dict/words | tail -1
```

3. Let `head` read from its standard input, but connect the standard input to a file:

```
head -23 < /usr/share/dict/words | tail -1
```

The < operator tells the UNIX command interpreter to connect the file `/usr/share/dict/words` to the input of `head`. The system tricks `head` into thinking its reading from the keyboard, but the input comes from the file – without requiring any change in the program!

4. Use the `cat` program to achieve the same effect as the preceding method. The `cat` program copies its input file(s) to the output. This extended pipeline therefore has the same effect as the preceding method:

```
cat /usr/share/dict/words | head -23 | tail -1
```

This longer pipeline is slightly less efficient than using the redirection operator. The pipeline requires an extra program (`cat`) copying its input to its output, whereas the redirection operator lets the lower level of the UNIX system achieve the same effect (replumbing the input) without the gratuitous copy.

As practice, let's use the UNIX approach to divide and conquer a search problem:

▶ *Imagine a dictionary of English alphabetized from right to left instead of the usual left to right. In other words, the dictionary begins with words that end in 'a'. In that dictionary, what word immediately follows* `trivia`*?*

This whimsical problem is drawn from a scavenger hunt [26]created by the computer scientist Donald Knuth, whose many accomplishments include the TEX typesetting system used to produce this book.

How is this different from 2? Just that you're "tricking it"? Is this in any way better than 2 or more useful in some other circumstance than 2?

By the same token, does it matter that the program now thinks its drawing from the keyboard input? Why would drawing from a pre-written file be any different for the computer, once you specify the file location?

I agree with these comments - I fail to see how this option is useful. Why would you ever want to "trick" the keyboard?

Sorry by keyboard I meant trick the system into believing it is receiving input from the keyboard.

I think it allows you to type in the whole line head -23<... the first time and from then on you can just type head -23 and leave out the file path. From then on the computer will assume that it should print from that path anyway.

I think, more importantly, the fact that we're being distracted by this question means that some of the detail here may be unnecessary to make the overall point, and seems to actually detract from our understanding.

I don't really understand what this is saying.

A minor point, but this should be "it's"

therefore reads lines typed at the keyboard, prints the 23rd line, and exits (even if the user is still typing).

2. Tell `head` to read its input from a file – for example from an English dictionary. On my GNU/Linux computer, the English dictionary is the file `/usr/share/dict/words`. It contains one word per line, so the following pipeline prints the 23rd word from the dictionary:

```
head -23 /usr/share/dict/words | tail -1
```

3. Let `head` read from its standard input, but connect the standard input to a file:

```
head -23 < /usr/share/dict/words | tail -1
```

The < operator tells the UNIX command interpreter to connect the file `/usr/share/dict/words` to the input of `head`. The system tricks `head` into thinking its reading from the keyboard, but the input comes from the file – without requiring any change in the program!

4. Use the `cat` program to achieve the same effect as the preceding method. The `cat` program copies its input file(s) to the output. This extended pipeline therefore has the same effect as the preceding method:

```
cat /usr/share/dict/words | head -23 | tail -1
```

This longer pipeline is slightly less efficient than using the redirection operator. The pipeline requires an extra program (`cat`) copying its input to its output, whereas the redirection operator lets the lower level of the UNIX system achieve the same effect (replumbing the input) without the gratuitous copy.

As practice, let's use the UNIX approach to divide and conquer a search problem:

Imagine a dictionary of English alphabetized from right to left instead of the usual left to right. In other words, the dictionary begins with words that end in 'a'. In that dictionary, what word immediately follows `trivia`*?*

This whimsical problem is drawn from a scavenger hunt [26]created by the computer scientist Donald Knuth, whose many accomplishments include the TEX typesetting system used to produce this book.

what is the point of all of this? the text should be more focused on the "estimation" techniques rather then get bogged down in the programming.

I think the section is trying to drive home the idea that unix programs illustrate divide&conquer techniques, but I agree it seems too concerned with the details of these programs.

I also agree. It's good to see how these topics are used in real examples, but this seems a little much.

I also agree with this. I am finding myself just trying to understand how the programs work with the filters and such and forgetting I am even supposed to think about divide and conquer.

Yeah. The relation seems fuzzy and unclear. I think the simplicity of divide and conquer needs to be used in this example. It seems way to focused on details. Details that drive the reader away from the original intent of the example

I agree. I actually learned some useful things about Unix from here, but that's not what we're here for.

what do you mean by redirection operator, do you refer to < ?

I think that just means you're reading the file from standard input rather than straight from the file.

Yes, you are right, but it would be nice if < was called that above when it is first mentioned.

In what way is this slightly less efficient? I see that it probably takes longer but is that where the inefficiency lies?

Does this actually create a phantom copy?

ooh this sounds cool. Is it weird that I'm excited to read on to figure out how to solve this?

This example helps to clarify some of the ideas somewhat, but this whole section is still largely incomprehensible to me.

Thanks for the example. At first, I thought this might be from z to a, but now I get that its based on last letter of the words.

therefore reads lines typed at the keyboard, prints the 23rd line, and exits (even if the user is still typing).

2. Tell `head` to read its input from a file – for example from an English dictionary. On my GNU/Linux computer, the English dictionary is the file `/usr/share/dict/words`. It contains one word per line, so the following pipeline prints the 23rd word from the dictionary:

```
head -23 /usr/share/dict/words | tail -1
```

3. Let `head` read from its standard input, but connect the standard input to a file:

```
head -23 < /usr/share/dict/words | tail -1
```

   The < operator tells the UNIX command interpreter to connect the file `/usr/share/dict/words` to the input of `head`. The system tricks `head` into thinking its reading from the keyboard, but the input comes from the file – without requiring any change in the program!

4. Use the `cat` program to achieve the same effect as the preceding method. The `cat` program copies its input file(s) to the output. This extended pipeline therefore has the same effect as the preceding method:

```
cat /usr/share/dict/words | head -23 | tail -1
```

   This longer pipeline is slightly less efficient than using the redirection operator. The pipeline requires an extra program (`cat`) copying its input to its output, whereas the redirection operator lets the lower level of the UNIX system achieve the same effect (replumbing the input) without the gratuitous copy.

As practice, let's use the UNIX approach to divide and conquer a search problem:

▶ *Imagine a dictionary of English alphabetized from right to left instead of the usual left to right. In other words, the dictionary begins with words that end in 'a'. In that dictionary, what word immediately follows* `trivia`*?*

This whimsical problem is drawn from a scavenger hunt [26]created by the computer scientist Donald Knuth, whose many accomplishments include the TEX typesetting system used to produce this book.

what is this

It's a 'typesetting' program. You may have heard of LaTeX before; it's commonly used in the scientific environment to format and print formulas which would otherwise be cumbersome to use in word processing.

how is this different from other typing programs?

Yeah, I would like a short description of what this system does and why it's unique.

googling TeX will explain the difference:

http://latexforhumans.wordpress.com/2008/10/07/why-use-latex/

or

http://ricardo.ecn.wfu.edu/ cottrell/wp.html

latex is really great and i love it. however, i don't think the details of how it works re needed here, in this book. it's more of a point in passing, to show how awesome donald knuth is.

I agree. It's just a little aside to embellish the achievements of Mr. Knuth. It doesn't really matter what TeX is (one can always Google it) for this material, and as Unix Rule #1 says "Small is beautiful". No need to clutter the book with random stuff.

The UNIX approach divides the problem into two parts:

1. Make a dictionary alphabetized from right to left.
2. Print the line following 'trivia'.

The first problem subdivides into three parts:

1. Reverse each line of a regular dictionary.
2. Alphabetize (sort) the reversed dictionary.
3. Reverse each line to undo the effect of step 1.

The second part is solved by the UNIX utility `sort`. For the first and third parts, perhaps a solution is provided by an item in UNIX toolbox. However, it would take a long time to thumb through the toolbox hoping to get lucky: My computer tells me that it has over 8000 system programs.

Fortunately, the UNIX utility `man` does the work for us. `man` with the `-k` option, with the 'k' standing for keyword, lists programs with a specified keyword in their name or one-line description. On my laptop, `man -k reverse` says:

```
  $ man -k reverse
  col (1)              - filter reverse line feeds from in-
put
  git-rev-list (1)     - Lists commit objects in reverse chrono-
logical order
  rev (1)              - reverse lines of a file or files
  tac (1)              - concatenate and print files in re-
verse
  xxd (1)              - make a hexdump or do the reverse.
```

Understanding the free-form English text in the one-line descriptions is not a strength of current computers, so I leaf through this list by hand – but it contains only five items rather than 8000. Looking at the list, I spot `rev` as a filter that reverses each line of its input.

▶ *How do you use* `rev` *and* `sort` *to alphabetize the dictionary from right to left?*

Therefore the following pipeline alphabetizes the dictionary from right to left:

Why/how does reverse alphabetizing the dictionary make this program better/more efficient?

I don't think this has to do with efficiency but is part of the problem definition/solution.

isn't this not really applicable to the divide and conquer idea- because its too hard to do this and where is the divide part

If I were asked to solve this problem on my own, I don't think I would come up with this as something I needed to do. I think it would help to explain why this is necessary.

it seems more natural to me to have these simply be the first 3 steps, for 4 steps total.

It is clear to me here that divide and conquer is being used in the problem but I think it should be stated at some point, probably earlier then this. This is the first time I realized we were breaking up the programming the same way we do problems.

Or maybe drawing a divide and conquer tree would be a useful way to show that?

Don't we want to get the next word after 'trivia' ending in -a in alphabetical order? But this seems like it would find the alphabetical ordered word after 'aivirt'

This was really helpful in that I could visualize the tree before we went through the problem, it would have also helped to have seen the second problem's 2 divisions verbally explained when it's branch was approached.

I understood what was being said here, but it was actually quite confusing the first time. Perhaps "reverse each word of a regular dictionary" would be more clear?

I agree with the replacing line with word, or simple clarifying we are talking about a unix dictionary where each line is a word.

I actually think this is a very helpful way to explain it for us to understand it.

This is sounding very much like an algorithmic breakdown. Interesting how algorithmic principle still prove useful in approximations.

I think it would be ususel in the class to go over some sorting algorithms, which was the first think that I think about with divide and conquer (merge sort etc) It's good for general knowledge and interview questions as well

The UNIX approach divides the problem into two parts:

1. Make a dictionary alphabetized from right to left.
2. Print the line following 'trivia'.

The first problem subdivides into three parts:

1. Reverse each line of a regular dictionary.
2. Alphabetize (sort) the reversed dictionary.
3. Reverse each line to undo the effect of step 1.

The second part is solved by the UNIX utility `sort`. For the first and third parts, perhaps a solution is provided by an item in UNIX toolbox. However, it would take a long time to thumb through the toolbox hoping to get lucky: My computer tells me that it has over 8000 system programs.

Fortunately, the UNIX utility `man` does the work for us. `man` with the `-k` option, with the 'k' standing for keyword, lists programs with a specified keyword in their name or one-line description. On my laptop, `man -k reverse` says:

```
  $ man -k reverse
  col (1)              - filter reverse line feeds from in-
put
  git-rev-list (1)     - Lists commit objects in reverse chrono-
logical order
  rev (1)              - reverse lines of a file or files
  tac (1)              - concatenate and print files in re-
verse
  xxd (1)              - make a hexdump or do the reverse.
```

Understanding the free-form English text in the one-line descriptions is not a strength of current computers, so I leaf through this list by hand – but it contains only five items rather than 8000. Looking at the list, I spot `rev` as a filter that reverses each line of its input.

▶ *How do you use* `rev` *and* `sort` *to alphabetize the dictionary from right to left?*

Therefore the following pipeline alphabetizes the dictionary from right to left:

What is meant by this? Is this just saying that there is one command that might do this for us but for the purposes of this argument we are using divide and conquer?

I know its in a slightly different font, but it's still hard to notice that 'man' and earlier 'cat' are actual functions

isnt this the same thing as going through the toolbox looking for an operation? you just know that it exists instead of looking for it.

it would be helpful if you can actually provide some examples here

This is a very useful section.

Yeah that would make this easier to follow.

Like many people have pointed out, I feel like I understand the outline given at the top of the page, and how that uses the divide and conquer method, but I'm not experienced with coding or UNIX, so all this seems very foreign to me, and the lack of examples makes it hard to follow

I feel the same way. I understand the idea behind divide-and-conquer, the example goes from a very understandable one (printing the 23rd line) to this, and I had trouble following the code and UNIX commands. I agree that more examples that build up to this problem would be beneficial in grasping the material.

This is the sort of question that would benefit from expanding the tree to solve it.

I'd say reverse each line. then sort. then reverse again

```
rev < /usr/share/dict/words | sort | rev
```

The second problem – finding the line after 'trivia' – is a task for the pattern-searching utility grep. If you had not known about grep, you might find it by asking the system for help with `man -k pattern`. Among the short list is

```
grep (1)              - print lines matching a pattern
```

In its simplest usage, grep prints every input line that matches a specified pattern. For example,

```
grep 'trivia' < /usr/share/dict/words
```

prints all lines that contain `trivia`. Besides `trivia` itself, the output includes `trivial`, `nontrivial`, `trivializes`, and similar words. To require that the word match `trivia` with no characters before or after it, give grep this pattern:

```
grep '^trivia$' < /usr/share/dict/words
```

The patterns are regular expressions. Their syntax can become arcane but their important features are simple. The ˆ character matches the beginning of the line, and the $ character matches the end of the line. So the pattern ˆtrivia$ selects only lines that contain exactly the text `trivia`.

▶ *This invocation of* `grep`*, with the special characters anchoring the beginning and ending of the lines, simply prints the word that I specified. How could such an invocation be useful?*

That invocation of `grep` tells us only that `trivia` is in the dictionary. So it is useful for checking spelling – the solution to a problem, but not to our problem of finding the word that follows `trivia`. However, Invoked with the `-A` option, `grep` prints lines following each matching line. For example,

```
grep -A 3 '^trivia$' < /usr/share/dict/words
```

will print 'trivia' and the three lines (words) that follow it.

When these pipelines are added into the text, it would be helpful if each word and operation was pointed out and explaining the transition and what gets done.

Again, why exactly do we need this? We could (from what I understood earlier) get the same output without using it, only the program would terminate at the end.

nice I was right

Is the 3 by itself is enough to have it print the 3 lines?

I've never been a unix user but after reading this chapter I am inclined to start, the pipeline system and general intuitive open source options are incredibly appealing.

Does it really make sense to have stored these now oddly-ordered words in the dictionary on the computer, rather than in a new file?

I don't understand why we are learning all this UNIX terminology. Don't we just want the concepts?

good to know but it seems like you're using this when you already know what you're looking for. Maybe you just use this to check it's there. Why not just say like open trivia or whatever you're trying to do since you know the exact name

Is this the loose equivalent of quotation marks around a phrase in a google search?

Yeah, but with the added requirement of 'trivia' being an entire line by itself.

What would of happened if there were no words similar to trivia that were next to it, it seems to me that this method would not have worked. Or maybe I don't understand what is going on here.

```
rev < /usr/share/dict/words | sort | rev
```

The second problem – finding the line after 'trivia' – is a task for the pattern-searching utility `grep`. If you had not known about `grep`, you might find it by asking the system for help with `man -k pattern`. Among the short list is

```
grep (1)             - print lines matching a pattern
```

In its simplest usage, `grep` prints every input line that matches a specified pattern. For example,

```
grep 'trivia' < /usr/share/dict/words
```

prints all lines that contain `trivia`. Besides `trivia` itself, the output includes `trivial`, `nontrivial`, `trivializes`, and similar words. To require that the word match `trivia` with no characters before or after it, give grep this pattern:

```
grep '^trivia$' < /usr/share/dict/words
```

The patterns are regular expressions. Their syntax can become arcane but their important features are simple. The ^ character matches the beginning of the line, and the $ character matches the end of the line. So the pattern t̂rivia$ selects only lines that contain exactly the text `trivia`.

This invocation of `grep`*, with the special characters anchoring the beginning and ending of the lines, simply prints the word that I specified. How could such an invocation be useful?*

That invocation of `grep` tells us only that `trivia` is in the dictionary. So it is useful for checking spelling – the solution to a problem, but not to our problem of finding the word that follows `trivia`. However, Invoked with the `-A` option, `grep` prints lines following each matching line. For example,

```
grep -A 3 '^trivia$' < /usr/share/dict/words
```

will print 'trivia' and the three lines (words) that follow it.

This attribute of grep seems almost too handy- is there no way if getting the position of trivia within this list of words?

It seems that those unfamiliar with UNIX would have a tough time figuring out all these commands - is there another way to do this?

I agree–the grep method and this -A option seem to basically perform the task for us, just in this one seemingly per-built in method(?). So aren't we just solving our problem using a method someone else already wrote to solve this problem?

I don't know UNIX and cannot understand if the -A option has to do with the fact that trivia has an A as the first letter when written backwards or if it is a general operation.

```
trivia
trivial
trivialities
triviality
```

To print only the word after 'trivia' but not 'trivia' itself, use `tail`:

```
grep -A 1 '^trivia$' < /usr/share/dict/words | tail -1
```

These small solutions combine to solve the scavenger-hunt problem:

```
rev </usr/share/dict/words | sort | rev | grep -A 1 '^trivia$'
| tail -1
```

Try it on a local UNIX or GNU/Linux system. How well does it work?

Alas, on my system, the pipeline fails with the error

```
rev: stdin: Invalid or incomplete multibyte or wide char-
acter
```

The `rev` program is complaining that it does not understand a character in the dictionary. `rev` is from the old, ASCII-only days of UNIX, when each character was limited to one byte; the dictionary, however, is a modern one and includes Unicode characters to represent the accented letters prevalent in European languages.

To solve this unexpected problem, I clean the dictionary before passing it to `rev`. The cleaning program is again the filter `grep` told to allow through only pure ASCII lines. The following command filters the dictionary to contain words made only of unaccented, lowercase letters.

```
grep '^[a-z]*$' < /usr/share/dict/words
```

This pattern uses the most important features of the regular-expression language. The ^ and $ characters have been explained in the preceding examples. The `[a-z]` notation means 'match any character in the range a to z – i.e. match any lowercase letter.' The `*` character means 'match zero or more occurrences of the preceding regular expression'. So ^[a-z]*$ matches any line that contains only lowercase letters – no Unicode characters allowed.

The full pipeline is

Comments on page 7

wait but aren't we sorting by the last letter? how is this helping us?

Have you not rearranged it yet by this point? Shouldn't all the words near trivia end in a if its reverse-alphabetized?

what is the point of this- I do not understand what I am supposed to get out of this that is useful

your previous line of code above had a "3" here instead of a "1", what just happened?

> the previous 3 was so he could see the next 3 words. Now that he just wants 1 word he replaces the 3 with a 1.
>
> > then why would he write | tail -1 also? isn't that redundant?
> >
> > > Because the grep would return "trivia, alluvia", and he only wants 'alluvia'
> > >
> > > Because the grep would return "trivia, alluvia", and he only wants 'alluvia'
> > >
> > > > very helpful, thank you!

I am having trouble following this code example.

Ok, so you are NOT actually saving your newly sorted list outside of the program, so to find the one after some other word (such as "dog") you would have to re-run the program

This is pretty easy to follow as I am doing it. This note also clarifies the confusion I had when I was trying to do the homework before reading this.

what does this means?

I don't really see why it matters what the computer says. It seemed as though everything before this worked and the extra steps taken don't really have to do with divide and conquer. Maybe all the technical terms are getting to me...

Again, why is this stuff relevant to this class?

where are these programs coming from?

so with one line of code you fixed the problem?

```
trivia
trivial
trivialities
triviality
```

To print only the word after ‘trivia’ but not ‘trivia’ itself, use `tail`:

```
grep -A 1 '^trivia$' < /usr/share/dict/words | tail -1
```

These small solutions combine to solve the scavenger-hunt problem:

```
rev </usr/share/dict/words | sort | rev | grep -A 1 '^trivia$'
| tail -1
```

Try it on a local UNIX or GNU/Linux system. How well does it work?

Alas, on my system, the pipeline fails with the error

```
rev: stdin: Invalid or incomplete multibyte or wide char-
acter
```

The `rev` program is complaining that it does not understand a character in the dictionary. `rev` is from the old, ASCII-only days of UNIX, when each character was limited to one byte; the dictionary, however, is a modern one and includes Unicode characters to represent the accented letters prevalent in European languages.

To solve this unexpected problem, I clean the dictionary before passing it to `rev`. The cleaning program is again the filter `grep` told to allow through only pure ASCII lines. The following command filters the dictionary to contain words made only of unaccented, lowercase letters.

```
grep '^[a-z]*$' < /usr/share/dict/words
```

This pattern uses the most important features of the regular-expression language. The ^ and $ characters have been explained in the preceding examples. The `[a-z]` notation means ‘match any character in the range a to z – i.e. match any lowercase letter.’ The * character means ‘match zero or more occurrences of the preceding regular expression’. So ^[a-z]*$ matches any line that contains only lowercase letters – no Unicode characters allowed.

The full pipeline is

explanation would be helpful here...not just another computer jargon

So any words with accents would be deleted? Wouldn’t this potentially change the solution if the word after trivia happened to have accents or uppercase letters

Yeah I’m not too happy with this compromise either, but I’m assuming there’s no other simple way to do it...

Some of these code examples are hard to follow.

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

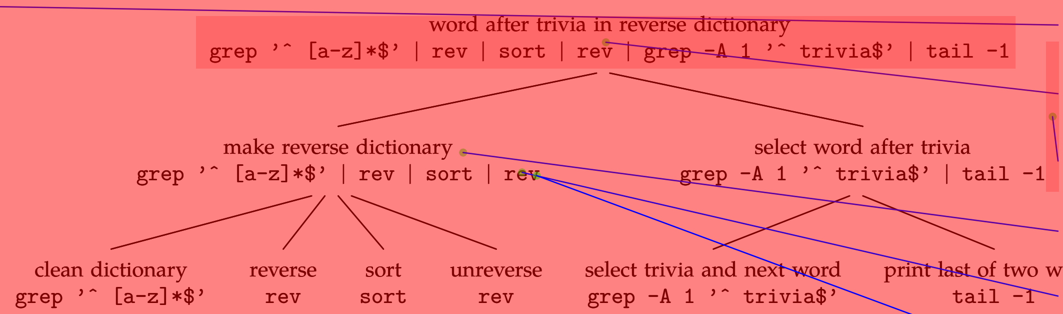

Running the pipeline produces produces 'alluvia'.

Problem 1.6 Angry
In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

I think by this stage in the chapter I got a bit lost...usually seeing a bunch of code frightens people. The chapter should go a little slower and maybe try and use simple pseudocode examples instead of focusing on UNIX.

I think I could've gotten the same out of the chapter without knowing the exact code. Could've used more reader friendly names for methods, but I do enjoy that I learned the something real

This is pretty simple to understand.

is this the code that reverses the dictionary?

tree takes a bit to understand, but makes sense

These trees with code in them are very difficult for me to understand.

this reminds me a lot of 6.001

This tree is hard to follow by itself.

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

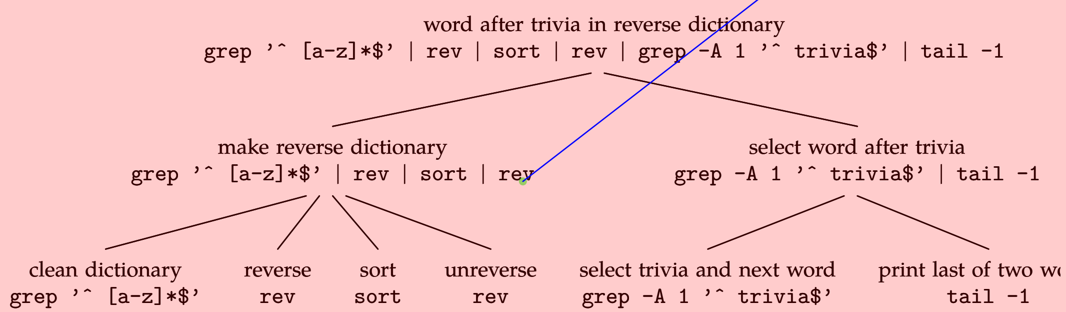

Running the pipeline produces produces 'alluvia'.

Problem 1.6 Angry
In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

I really like this diagram of a tree explaining the whole pipeline. It really ties back what we know of trees and divide and conquer to the UNIX example.

Although this tree certainly helps, a lot of this reading was very confusing to me as someone with extremely limited coding experience. I was able to follow the UNIX principles and even the basic head and tail topics, but after that I was pretty lost. Using this tree I was able to go back and understand a little bit more, but I am certainly not completely comfortable with the material.

And I guess to voice the opposite view point, this was an extremely enlightening example as it provided an interesting way for me to consider a problem I'd otherwise think of in a very informal fashion if I had to solve it, and probably take a lot longer to figure out, even with the experiences using Unix/Linux I have. So Thanks!

It would have helped me to see this, or at least parts of this, as we went through the problem and not the whole thing at the end.

I totally agree. Until I saw this tree I had totally lost the divide and conquer to trying to understand UNIX.

Although placing the tree at the end was helpful it would have been easier to understand if less UNIX coding was used and more english and logic was used.

Same here, I am not very familiar with UNIX but this tree really helped me figure out what was going on. It may be helpful to show this in the beginning as well

Maybe the best solution is instead to show the tree being developed as it progresses. Like when we realize we have to clean the dictionary, we can see that part added to the tree - the problem as described at the beginning of the section has a simpler tree that we add to so why not show the initial state and the complications added to it?

I agree, starting with a basic tree which was then improved upon would keep me focused on the divide and conquer concept as opposed to the syntax of UNIX

I strongly second the recommendation to slowly build the tree. I'm a relative noob in programming, so it was a bit difficult to follow the functions (since it's an entirely new syntax to me). However, if you could show the tree as it went along and then have it here as a cumulative one, that would be awesome.

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

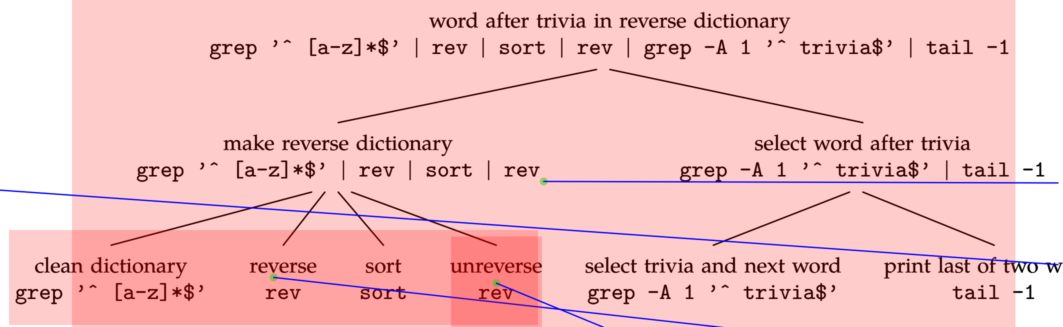

Running the pipeline produces produces 'alluvia'.

Problem 1.6 Angry
In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

I agree. This tree is very helpful. Some of the above text is hard to follow, but this clears it up nicely.

Yeah I definitely agree too. I thought the sections about were long and somewhat difficult to follow, if someone doesn't have very much UNIX experience. This tree diagram, is a great way to represent the previous steps visually and to tie it back to the divide&conquer method we learned about in the previous sections.

I also like this tree, and feel that perhaps it(or a smaller version) would be helpful to have it earlier, because I began to jot a few words down in order to follow the logic

Yeah this really sums up everything well. While I don't quite understand all the code, this tree has helped me understand the divide and conquer application.

Yeah this is super helpful. It clarifies the whole process.

This tree really helps understand how the problem is solved- maybe it would have been better to put earlier since this explanation is so long

these aren't independent calculations. each one should be performed after the other, so they shouldnt be represented in this leaf format.

Aye, I agree. Are we supposed to assume that we apply these steps from left to right?

I'm confused about why this goes in the bottom part of the tree. Don't we want to select the word after trivia then unreverse it?

It would be helpful to see how exactly we got to alluvia from the tree.

I feel like it has been mostly described, but maybe seeing a short list of the two words proceeding and following trivia in the reverse dictionary would give a nice visual to wrap up the problem.

We would need to write a program to answer this, right? If so, then how does approximation come into this?

Can a person who is unfamiliar with Unix answer this question? I have no idea how I would begin (though I do understand the previous example).

...or any form of problem solving

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

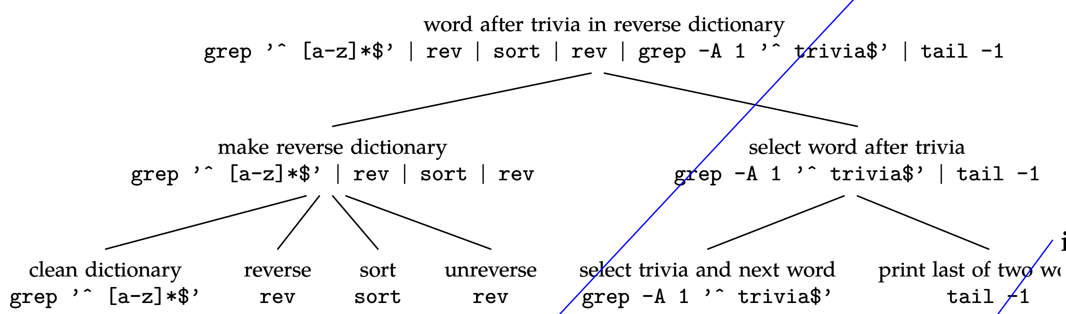

Running the pipeline produces produces 'alluvia'.

Problem 1.6 Angry
In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

I feel like I understood divide and conquer okay before this chapter, but as someone who doesn't know anything about UNIX, this chapter was a hard read and I didn't get much out of it in the end...

I think he's just trying to make the point that divide and conquer is a common, useful tool for solving problems even when used outside of the confines of estimation. However, it would probably help if this was stated at the beginning of the chapter.

I agree, this helped me to realize that divide and conquer is used in other instances but the technical pieces of UNIX are unfamiliar to me and made the section harder to understand.

I also agree. I think divide and conque is something people generally do, especially when things are exact like computer code. I'd like more talk about how to go about breaking things down. Hints, tips, etc..

it wouldve been better to leave it short.

Agreed, I got so lost in the programming discussion that when we came back to divide and conquer I felt like I had missed the whole point.

Definitely agreed. I am even in course 6 and I understood all the details about UNIX and programming, and I still thought this example was excessive. There was way too much explanation necessary for such a simple concept as divide and conquer. And no where throughout this example did I even remember that this was supposed to be about divide and conquer. I would delete this whole example from the book.

I think we can find a better example to use to show how common divide-and-conquer reasoning is.

A very good example of divide and conquer, but the result is very discrete and precise: not an approximation or an estimation.

I think the long explanation of various unix functionalities didn't add to the explanation and instead made it all significantly more confusing, even though I understand how to use most of them.

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

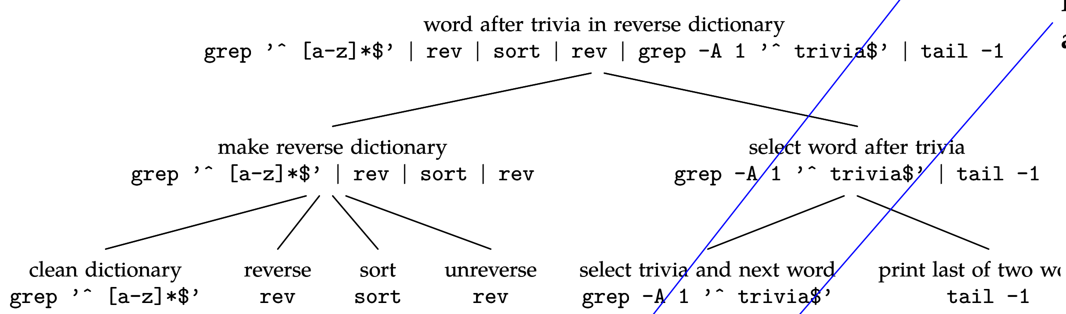

Running the pipeline produces produces 'alluvia'.

Problem 1.6 Angry
In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

The UNIX Part of this doesn't quite help the final message; I feel like it's rather off topic from Divide and Conquer. It's cool to learn, and I understand what it's trying to prove, but I just don't get the helpfulness to divide and conquer. Maybe after lecture it will become more clear.

Is the UNIX example relevant specifically? Would other programs be just as suitable as an example as long as they are modular in design?

I'm still not sure I understand until where I expand the tree. Is there a rule of thumb about where to stop expanding?

I agree. In this case it looks like the author expanded the tree until he could use single methods to solve each problem (without combining methods with the |). But I'm not experienced with UNIX or coding so I'm not sure...

You should keep going until you have numbers that are manageable and are likely to be correct (at least to the order of magnitude) or numbers that you know are correct. Typically, these numbers are much smaller than the actual answer. In a way, we're approximating a big thing with many small things that combine together.

More generally, expand until each problem (leaf) is soluble on its own. This may vary for different people. Just think of how you carry out any task that involves more than one step or piece.

was this explicitly stated before? it makes sense though.

This is true for most divide and conquer cases where you're trying to estimate but it doesn't seem applicable here. Breaking down the problem makes it easier to realize how to code for an answer, but since you're relying on a computer program to find the answer, accuracy doesn't seem to be a concern. This point was much clearer in the economics example.

I really think divide and conquer is the easiest way to solve complicated problems. Also, knowing how to break the problem into steps is extremely helpful, but I find that to be the most difficult part.

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

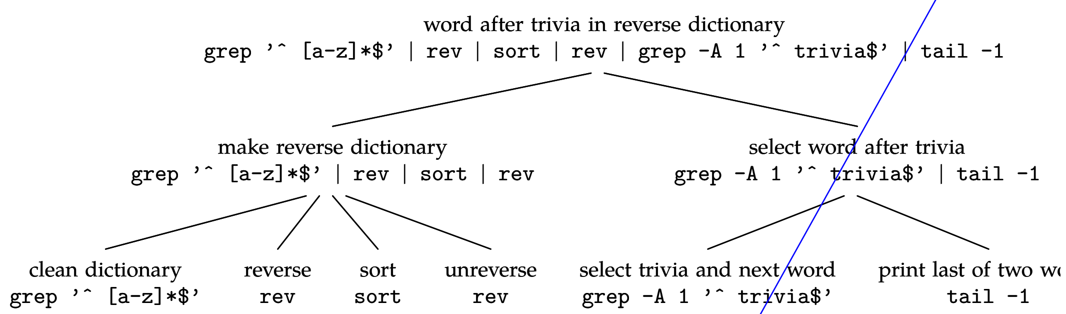

Running the pipeline produces produces 'alluvia'.

Problem 1.6 Angry
In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

does this actually ensure accuracy? it seems like the more you subdivide, the more you need to multiply guessed numbers together, possibly getting a much larger answer than expected.

I know someone brought this up on the earlier section about trees, but does the error multiply similarly up the tree? What is the relationship between the size of the error and the number of leaves? (I doubt its a constant for all types of problems but it would still be interesting to study)

I feel like this is just a reflexive fear. Everybody seems to keep fearing the possibility of errors compounding and snowballing out of control, but from something as simple as the health cost example in class, I've found that if you keep trying approaches to a problem until you reach an angle where you're comfortable with the numbers, the numbers won't be too radically off.

Just an observation.

Well I like the point is to break the problem down into things you can estimate with accuracy and get rid of that error.

It seems that trees also represent nearly all other kinds of problem solving and logical processes. Engineering system design, analytical thinking, etc.

That's an interesting thought. I agree it is pretty useful, but I'm wondering about the efficiency of using trees over other methods.

There are other diagramming options for other types of problems like flow charts, organizational charts, Gantt... Trees are useful but not universally, as we've seen with our difficulties in illustrating connections and redundancy.

And in books, like this chapter... I'd still prefer something different than the text processing example, but maybe that's just me.

I think it's a good idea to have some example that's not computing numbers, although it doesn't have to be a text processing. But it might be this is the simplest example to read without Linux experience and potentially grasp the first time.

```
grep '^[a-z]*$' < /usr/share/dict/words \
| rev | sort | rev \
| grep -A 1 '^trivia$' | tail -1
```

where the backslashes at the end of the lines tell the shell to continue reading the command beyond the end of that line.

The tree representing this solution is

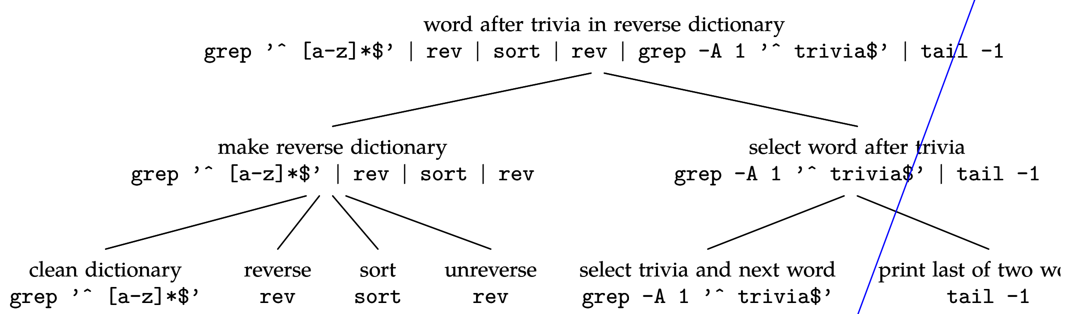

Running the pipeline produces produces 'alluvia'.

> **Problem 1.6 Angry**
> In the reverse-alphabetized dictionary, what word follows angry?

Although solving this problem won't save the world, it illustrates how divide-and-conquer reasoning is built into the design of UNIX. In short, divide and conquer is a ubiquitous tool useful for estimating difficult quantities or for designing large, successful systems.

Main messages

This chapter has tried to illustrate these messages:

1. Divide large, difficult problems into smaller, easier ones.
2. Accuracy comes from subdividing until you reach problems about which you know more or can easily solve.
3. Trees compactly represent divide-and-conquer reasoning.
4. Divide-and-conquer reasoning is a cross-domain tool, useful in text processing, engineering estimates, and even economics.

Could you possibly provide an example of how divide-and-conquer reasoning is used in economics? I think it would be helpful for future reading of this section, but I would also be interested in learning about an example currently.

I agree, an example would be nice. Otherwise this point feels I little out of place in the conclusion, given we haven't seen such an example before.

The oil problem is sort of economics...ish. More so than the others.

It seems that by saying "and even economics," that divide and conquer is used for only the problems mentioned, but I think you mean to say that it applies to many problems in life, sometime obvious and sometimes not.

I think by saying even economics, the author is trying to say divide-and-conquer's usefulness expands more than just the examples mentioned and suggests a possible field you hadn't thought of.

I don't think we need an example of that. Especially not at this part of the text. There are already some stellar examples of how this works and we don't actually need an economics specific one to udnerstand it better. This bullet point just serves to illustrate how wide-reaching this technique is beyond just estimation (plus it wouldn't really fit here... I mean, it's a bullet point).

By breaking hard problems into comprehensible units, the divide-and-conquer tool helps us organize complexity. The next chapter examines its cousin abstraction, another way to organize complexity.

Problem 1.7 Air mass
Estimate the mass of air in the 6.055J/2.038J classroom and explain your estimate with a tree. If you have not seen the classroom yet, then make more effort to come to lecture (!); meanwhile pictures of the classroom are linked from the course website.

Problem 1.8 747
Estimate the mass of a full 747 jumbo jet, explaining your estimate using a tree. Then compare with data online. We'll use this value later this semester for estimating the energy costs of flying.

Problem 1.9 Random walks and accuracy of divide and conquer
Use a coin, a random-number function (in whatever programming language you like), or a table of reasonably random numbers to do the following experiments or their equivalent.

The first experiment:

1. Flip a coin 25 times. For each heads move right one step; for each tails, move left one step. At the end of the 25 steps, record your position as a number between −25 and 25.
2. Repeat the above procedure four times (i.e. three more times), and mark your four ending positions on a number line.

The second experiment:

1. Flip a coin once. For heads, move right 25 steps; for tails, move left 25 steps.
2. Repeat the above procedure four times (i.e. three more times), and mark your four ending positions on a second number line.

Compare the marks on the two number lines, and explain the relation between this data and the model from lecture for why divide and conquer often reduces errors.

Problem 1.10 Fish tank
Estimate the mass of a typical home fish tank (filled with water and fish): a useful exercise before you help a friend move who has a fish tank.

Problem 1.11 Bandwidth
Estimate the bandwidth (bits/s) of a 747 crossing the Atlantic filled with CDROM's.

Comments on page 9

So the point of this was to help us organize big problems, not estimate. Ah.

copy edit: I think there should be a comma between cousin and abstraction, since cousin is the object, not an adjective. Abstraction is then in apposition to cousin.

Is there a constant for the density of air that we should know?

Is there a section of the book with useful constants? Maybe a page or a couple pages like the one in our desert island pretest? That would be a pretty nifty trivia cheat sheet.

OCW Spring 2008 for this class has a version of the book, and page 2 is Back-of-the-envelope numbers, basically the same as the desert island sheet. You can start making your own based on these, but please don't beat me at trivia with it.

For the record, will this book give solutions to these "try-for-yourself" type questions?

It should in the ideal world. But I suspect that, in order to actually finish the book, I may just have to release it into the world without them. But it will be a freely licensed book, and I hope that people will contribute solutions...

Who moves a full fish tank bigger than a few gallons?

Problem 1.12 Repainting MIT
Estimate the cost to repaint all indoor walls in the main MIT classroom buildings. [with thanks to D. Zurovcik]

Problem 1.13 Explain a UNIX pipeline
What does this pipeline do?

```
ls -t | head | tac
```

[Hint: If you are not familiar with UNIX commands, use the `man` command on any handy UNIX or GNU/Linux system.]

Problem 1.14 Design a UNIX pipeline
Make a pipeline that prints the ten most common words in the input stream, along with how many times each word occurs. They should be printed in order from the the most frequent to the less frequent words. [Hint: First **tr**anslate any non-alphabetic character into a newline. Useful utilities include `tr` and `uniq`.]

Nice, I really like this question, but it should be more general when it's published.

My UNIX underlying OS X does not have a tac man entry or allow me to run that program.

2
Abstraction

2.1	Diagrams	28
2.2	UNIX abstractions	31
2.3	Recursion	35
2.4	Low-pass filters	38
2.5	Summary and further problems	38

Divide-and-conquer reasoning breaks enigmas into manageable problems. When the reasoning is represented as a tree, the manageable problems become the leaf nodes of the tree, and they are conceptually simpler than the original problem or its intermediate subproblems. For example, the length of a classical symphony is a simple concept compared to the data capacity of a CDROM.

Being simpler, it is more likely than the parent nodes to be used in another calculation. Imagine that you are an architect designing a classical concert hall. One task is to ensure sufficient airflow to handle the heat produced by 1500 audience members during a concert. But how long is a concert? Reuse the symphony leaf node from the CDROM-capacity estimate. Concerts often include a symphony before or after a break (the intermission), with a comparably long other half, so a rough concert duration 2.5 hours.

Creating and using such reusable parts is the purpose of our second tool for organizing complexity: abstraction. Abstraction is, according to the *Oxford English Dictionary* [29]:

> The act or process of separating in thought, *of considering a thing independently of its associations;* or a substance independently of its attributes; or an attribute or quality independently of the substance to which it belongs. [my italics]

The most important characteristic of abstraction is reusability. As Abelson and Sussman [1, s. 1.1.8] describe:

So after reading this section, and the comments, I think, but I'm not sure if I got this concept correctly (it's a bit confusing to read at first). Is the idea of abstraction to basically make "templates" in order to reduce/simplify/organize a task, and so that things lower on the "tower" are unimportant to continue? If that is the case, then this example makes sense finally at the end, by showing that a program/code designed to do a task does it better than just using a more general program that cannot do the specific tasks? I'm trying to think of other more, non-obvious real life uses of abstraction and am having a hard time thinking of a good one?

> A resistor of a specified resistance is an abstraction of what is really an analog transfer curve with tolerances. All the elements are abstractions. We can think about them without having to think about all their quarks and gluons or draw abstract diagrams like O-H-O. Scaling up, we can think about a bucket of water without imagining every molecule and trying to keep track of its position. Think of it as collapsing a lot of information into a more manageable concept. It's key to engineering so you probably do it all the time without thinking.

I think I'm having a lot of trouble seeing how this section fits in with approximations. It seems like the main point of the chapter is that you can write code to produce a tree easily. I also don't quite understand what abstraction in this context is.

This is a very vague, hand-waving definition of abstraction.

what is this boxes package? and how do we know which level of abstraction to look at without testing?

how did you come up with these lines? which level of abstraction are you looking at? I find this to be a very confusing example overall...I didn't realize all the scripts above were describing parts of an abstraction tower.

aren't most drawing programs designed this way? maybe I don't understand exactly what you mean by graphical captive UI, but this way the user has lots of options to format the tree as he/she wished.

So is abstraction using precreated shortcuts in creative ways? Also, I imagine it took a while to make the program. Doesn't it kinda cancel out?

2
Abstraction

2.1 Diagrams 28
2.2 UNIX abstractions 31
2.3 Recursion 35
2.4 Low-pass filters 38
2.5 Summary and further problems 38

Divide-and-conquer reasoning breaks enigmas into manageable problems. When the reasoning is represented as a tree, the manageable problems become the leaf nodes of the tree, and they are conceptually simpler than the original problem or its intermediate subproblems. For example, the length of a classical symphony is a simple concept compared to the data capacity of a CDROM.

Being simpler, it is more likely than the parent nodes to be used in another calculation. Imagine that you are an architect designing a classical concert hall. One task is to ensure sufficient airflow to handle the heat produced by 1500 audience members during a concert. But how long is a concert? Reuse the symphony leaf node from the CDROM-capacity estimate. Concerts often include a symphony before or after a break (the intermission), with a comparably long other half, so a rough concert duration 2.5 hours.

Creating and using such reusable parts is the purpose of our second tool for organizing complexity: abstraction. Abstraction is, according to the *Oxford English Dictionary* [29]:

> The act or process of separating in thought, *of considering a thing independently of its associations;* or a substance independently of its attributes; or an attribute or quality independently of the substance to which it belongs. [my italics]

The most important characteristic of abstraction is reusability. As Abelson and Sussman [1, s. 1.1.8] describe:

could we go over divide and conquer vs. abstraction conceptually? they overlap on some key parts and differ in some other key parts. I'd just like to see that line between them more distinctly

I still don't understand abstraction after reading this section...

Kind of confused...mostly because I've never done programming but does this seek to compare abstraction to reusing portions of the GUI computer program?

I'm very confused, can we have a non programming example to explain abstraction? I think I understand the point of abstraction but the programming portion has completely lost me.

Unlike the MH example in class, it seems that using abstraction to make a tree is longer and more time consuming than using a prescribed program. Is there an example where using abstraction is time efficient?

2 Abstraction

2.1 Diagrams 28
2.2 UNIX abstractions 31
2.3 Recursion 35
2.4 Low-pass filters 38
2.5 Summary and further problems 38

Divide-and-conquer reasoning breaks enigmas into manageable problems. When the reasoning is represented as a tree, the manageable problems become the leaf nodes of the tree, and they are conceptually simpler than the original problem or its intermediate subproblems. For example, the length of a classical symphony is a simple concept compared to the data capacity of a CDROM.

Being simpler, it is more likely than the parent nodes to be used in another calculation. Imagine that you are an architect designing a classical concert hall. One task is to ensure sufficient airflow to handle the heat produced by 1500 audience members during a concert. But how long is a concert? Reuse the symphony leaf node from the CDROM-capacity estimate. Concerts often include a symphony before or after a break (the intermission), with a comparably long other half, so a rough concert duration 2.5 hours.

Creating and using such reusable parts is the purpose of our second tool for organizing complexity: abstraction. Abstraction is, according to the *Oxford English Dictionary* [29]:

> The act or process of separating in thought, *of considering a thing independently of its associations;* or a substance independently of its attributes; or an attribute or quality independently of the substance to which it belongs. [my italics]

The most important characteristic of abstraction is reusability. As Abelson and Sussman [1, s. 1.1.8] describe:

Read this introduction to our next chapter (the reading is 4 pages) and submit the memo by 9am on Wednesday (2/17).

Isn't this intro still mostly about divide-and-conquer? Maybe it should go back in the first chapter?

I expected this to involve using diagrams to abstract, not using abstraction on the process of making diagrams. – I take this back in part because I didn't know the 2.1 section on diagrams hadn't started here.

These last few sections have been particularly CS-based; I know that many of the students have not enjoyed it as much.

These last few sections have been particularly CS-based; I know that many of the students have not enjoyed it as much.

If the window is left open, the new readings don't show up unless you hit refresh. It's almost tricked me a few times.

This is an absolutely lovely breakdown of divide-and-conquer; it's simple, concise, and covers the main points. It'd be helpful to see it earlier in the d-and-c section.

Can't an architect use divide and conquer for this question?

Yes, I think that a divide and conquer approach could be used for this. It does seem like a difficult calculation though.

strangely i don't ever remember being hot during a concert

I think that it's because it's a problem that has already been figured out. Concert halls have been around so long that not much thought has to be given about airflow.

2
Abstraction

2.1 Diagrams 28
2.2 UNIX abstractions 31
2.3 Recursion 35
2.4 Low-pass filters 38
2.5 Summary and further problems 38

Divide-and-conquer reasoning breaks enigmas into manageable problems. When the reasoning is represented as a tree, the manageable problems become the leaf nodes of the tree, and they are conceptually simpler than the original problem or its intermediate subproblems. For example, the length of a classical symphony is a simple concept compared to the data capacity of a CDROM.

Being simpler, it is more likely than the parent nodes to be used in another calculation. Imagine that you are an architect designing a classical concert hall. One task is to ensure sufficient airflow to handle the heat produced by 1500 audience members during a concert. But how long is a concert? Reuse the symphony leaf node from the CDROM-capacity estimate. Concerts often include a symphony before or after a break (the intermission), with a comparably long other half, so a rough concert duration 2.5 hours.

Creating and using such reusable parts is the purpose of our second tool for organizing complexity: abstraction. Abstraction is, according to the *Oxford English Dictionary* [29]:

> The act or process of separating in thought, *of considering a thing independently of its associations;* or a substance independently of its attributes; or an attribute or quality independently of the substance to which it belongs. [my italics]

The most important characteristic of abstraction is reusability. As Abelson and Sussman [1, s. 1.1.8] describe:

So, the ones we reuse a lot (in this class and life in general) are the ones we should memorize? Like the number of seconds in a year?

By using them often, you're also more likely to remember them rather than memorizing by rote.

Agreed - I think we'll just see them often enough that over time we'll remember them. Plus, I think we can raise flags in our minds that say "hey, this number sounds useful!" - seconds in a year, population of the US, etc should all fall into that category. In this case, where we talk about CD-ROM capacity, I'm sure it's fine to just look it up.

It would be nice if you finished solving the problem in this section.

i suppose it just depends on which values you are more familiar with.

I agree, I think it's important to emphasize that there is not one way to break down these problems, but instead it relies a lot on the pieces of information you individually know.

Maybe it will become more clear as I continue reading, but so far, I'm struggling to understand the difference between divide and conquer and abstraction.

This technique seems like it comprises divide and conquer, making trees, and robustness. I'm struggling to see what else it contains/employs besides these three things...?

How does this connect to divide and conquer? meaning treat the leaves as separate things?

I thought today's example about tracing the origin of species and spread of the bible was a great example of this

I've never been a fan of dropping dictionary definitions. But the idea of Abstraction is to generalize your concept, and I think this article just put in the only definition of abstraction that fits for their use.

I'm not sure that a dictionary definition is really the most useful definition of abstraction here. It might work better to just define it within the context of our goals towards estimation.

Yeah I feel like its kind of vague in terms of usefulness...

2
Abstraction

2.1	Diagrams	28
2.2	UNIX abstractions	31
2.3	Recursion	35
2.4	Low-pass filters	38
2.5	Summary and further problems	38

Divide-and-conquer reasoning breaks enigmas into manageable problems. When the reasoning is represented as a tree, the manageable problems become the leaf nodes of the tree, and they are conceptually simpler than the original problem or its intermediate subproblems. For example, the length of a classical symphony is a simple concept compared to the data capacity of a CDROM.

Being simpler, it is more likely than the parent nodes to be used in another calculation. Imagine that you are an architect designing a classical concert hall. One task is to ensure sufficient airflow to handle the heat produced by 1500 audience members during a concert. But how long is a concert? Reuse the symphony leaf node from the CDROM-capacity estimate. Concerts often include a symphony before or after a break (the intermission), with a comparably long other half, so a rough concert duration 2.5 hours.

Creating and using such reusable parts is the purpose of our second tool for organizing complexity: abstraction. Abstraction is, according to the *Oxford English Dictionary* [29]:

> The act or process of separating in thought, *of considering a thing independently of its associations;* or a substance independently of its attributes; or an attribute or quality independently of the substance to which it belongs. [my italics]

The most important characteristic of abstraction is reusability. As Abelson and Sussman [1, s. 1.1.8] describe:

I understand the idea of considering things independently- but isn't the whole point of divide and conquer to consider things as they relate to the goal? In order to divide and conquer, we must think about what smaller things contribute/ relate to the bigger picture so it's hard to then not consider them as parts of a whole.

I agree–it sounds a lot like divide and conquer, but I think abstraction goes deeper than divide and conquer. in divide and conquer we're looking into different components but in abstraction we're finding different ways of approaching a problem by seeing what it breaks down into...

I agree as well, it seems like abstraction is saying "step away and look again" rather than "how can we break this into bite sized chunks"

Agreed. I see there being somewhat of a balance- if things are too broad they aren't very useful, but if they are too specific they can't be reused. The trick is to find things that are small enough that they are manage-able and still be able to use the same numbers in a variety of applications

I agree that this is ambiguous - perhaps Sanjoy should use his own words instead of a dictionary definition to clarify confusion.

I also thought that the definition of abstraction kind of sounds like the same thing as divide and conquer. In my other classes abstraction is a way of viewing a variable at an acceptably high enough level. For example, when i learned about circuits, we "abstracted" away from analyzing it using really horrendous Maxwell equations, but "abstracted" components using values of Resistance, Inductance, and Capacitance. Those values of R,L, and C contain a lot of physics, but since we don't really care about the physics to solve the final solution, abstraction was used to lump unnecessary.

I don't really think this explains abstraction at all. I've read the entire document and I'm still not really sure what it is or how it helps in making approximations. It looks like it could be useful in making a diagram or a program, but how does that relate to approximating?

Is it possible to include an approximation-related example? Since it is the introduction, I see how the coding example can explain the concept. But it took me several readings to understand it.

Why is it helpful to think of something independently when it goes with something else to give you an answer.

2
Abstraction

2.1 Diagrams 28
2.2 UNIX abstractions 31
2.3 Recursion 35
2.4 Low-pass filters 38
2.5 Summary and further problems 38

Divide-and-conquer reasoning breaks enigmas into manageable problems. When the reasoning is represented as a tree, the manageable problems become the leaf nodes of the tree, and they are conceptually simpler than the original problem or its intermediate subproblems. For example, the length of a classical symphony is a simple concept compared to the data capacity of a CDROM.

Being simpler, it is more likely than the parent nodes to be used in another calculation. Imagine that you are an architect designing a classical concert hall. One task is to ensure sufficient airflow to handle the heat produced by 1500 audience members during a concert. But how long is a concert? Reuse the symphony leaf node from the CDROM-capacity estimate. Concerts often include a symphony before or after a break (the intermission), with a comparably long other half, so a rough concert duration 2.5 hours.

Creating and using such reusable parts is the purpose of our second tool for organizing complexity: abstraction. Abstraction is, according to the *Oxford English Dictionary* [29]:

> The act or process of separating in thought, *of considering a thing independently of its associations;* or a substance independently of its attributes; or an attribute or quality independently of the substance to which it belongs. [my italics]

The most important characteristic of abstraction is reusability. As Abelson and Sussman [1, s. 1.1.8] describe:

This is a great word to describe the use of abstraction. An abstracted element needs to be a contained entity. Something kind of like UNIX programs that can be reused for different tasks.

This ties in nicely with the modular part of the UNIX philosophy. Maybe this section should come before that one?

Or at least this section should briefly mention the connection to UNIX.

So this makes it sound like they want us to remember every estimation we ever made. I understand it's easier if you already estimated to just use it again, but I think it's pretty unpractical in reality.

The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

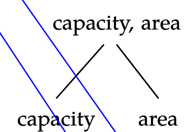

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

I am starting to get a little confused here, connecting the dots.

this really clarifies a lot

This is much more useful than the Oxford dictionary definition.

I agree especially when considering how we are using the word.

I agree as well. Is it really necessary to have the dictionary definition? I found the OED def to be a little broad given the course.

I also agree! The Oxford definition probably shouldn't be used and this put in its place.

i disagree. the first definition explains "abstraction" while this one is a more intuitive practice.

I like having the two definitions since the first one describes the general definition of abstraction while the second tells how it relates to what we are learning.

Well, this definition is the object oriented programming version specifically.

I also think that the oxford definition is kind of useless. I actually just skimmed over it because it seemed fairly complex and I assumed it would be better explained later.

I like the dictionary definition because it gives a technical introduction to what you're actually trying to accomplish. After knowing the process, the description gives a more thorough definition of the thought process.

I was thinking the same like. Methods pop right into mind.

So abstraction goes further than divide and conquer- it is a type of divide and conquer but goes further to generalize each solution

abstraction!!! It's pretty cool that divide and concur methods of approximation are analogous to methods in computer programing

The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

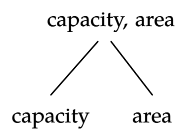

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

Choosing appropriate modules seems to be the critical and important step in this process. However, sometimes it can be very difficult to decide exactly how to make these modules. Are there any rules or guidelines to ensure that you modularize things in a way that is most helpful?

One way is to watch your own actions with slightly blurry vision. If you find yourself doing the same thing, or making the same calculation a second or third time, then you are likely to have a spot to make an abstraction.

The reason I say "slightly blurry vision" as well is that if you look too closely, then you'll never think you are doing something for a second time. Rather, you'll find all the differences. But with blurry vision, those details will hopefully disappear ("get abstracted away") and you'll be left with the common features.

i feel like this material is much more accessible to the 6.055 students than to the 2.038 students.

Yes... I'm a little nervous about this unit because I have no programming experience.

I think he's more interested in the ideas he's already explained than in the actual example. He's just trying to show how you can actually use stuff we've learned to solve problems.

I am not sure if that's true though, the text does go into detail...

i think the point of the statement "what they write about programs applies equally well to understanding other systems" is supposed to reassure us about the fact that we don't just have to be course 6 to understand and benefit from this.

Can we see an example of abstraction (or any other unit for that matter) that applies mech e concepts?

The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

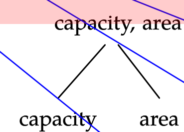

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

You mention this "tower" later on in the section and I think it really helped me understand this better; that things on the bottom of the tower become "abstracted" (quarks are not important to the understanding of fluids, but they still exist). I think if the tower analogy is defined more clearly this would be really helpful for a lot of people!

I agree. The appearance of the 'abstraction tower' here seems rather abrupt and caused me to look through previous paragraphs thinking that I had missed its definition already.

A figure for this example of an 'abstraction tower' would also be awesome.

so is a tower what's used for abstraction and a tree for divide and conquer?

I agree with the other comments here - the "abstraction tower" suddenly appears with no explanation whatsoever. A diagram or short description would be very helpful.

perhaps "abstract tower" (in quotes) is better? you would take it more metaphorically rather than expecting a definition. after all, everyone knows that a tower is.

This is a much better visualization of abstraction for more visual people.

why are we starting from the bottom here? are we going backwards from the divide and conquer method- I don't know why you would ever start with the smallest form of matter anyway- it seems to trivial to me (sorry I don't really like chemistry)

I like that you use such a simple example to follow the definition, but is there one that would be even more accessible. It helps for people to be able to visualize things and it was hard for me to visualize the relationships between the elements of this example because of my backgrounds.

I actually thought this was a very helpful example. It breaks down something as complicated as a fluid to suff as simple as quarks. (I am also taking some course 8 classes, so I love this analogy)

I REALLY liked this example. Maybe I am just a physics major, but I would like to point out it was one of my favorite examples yet on the reading.

Using "like" makes it sound as if there is merely an analogy between the behavior of molecules and the behavior of fluids. But large collections of molecules do not just act like a fluid (under the right conditions), they actually are a fluid.

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

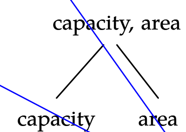

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

I don't understand this reference.

In this example, I see the abstraction tower. It's pretty clear to me the different levels of "detail". But I do not understand the tree language one below that well...

This was a helpful example; generic examples are always excellent for keeping a reader's attention.

so is an abstraction is an easy to work with model? Like how we use circuits to model conduction, or control systems, or eletromagnetic phenomena.

Yeah, and that's why the resistors, capacitors, and inductors we use are part of the "lumped matter" abstraction.

I feel that the example of the fluid is more relevant and intuitive than the concept of abstraction in programming

I feel like that's the point of abstraction. Quarks, and molecules would be useless to use in order to understand airflow but as they are looked at as a fluid these problems can be solved. I see abstraction more as a tool for creating intuitive solutions then the solution itself.

I agree with this abstraction is more of a method, or a means to an end.

This fluid example helped me wrap my head around the subject more.

this makes me want to yell: Abstraction! yippie

but you still need to understand fluid behavior in order to use it.

I guess this is the most direct contrast to the divide and conquer, where the point is to know exactly how they interact so you can do calculations

At the same time, we should recognize that the ends of abstraction is the same as divide-and-conquer: get onto terms with which you are comfortable generating numbers.

How can this help generate numbers f we are just examining big-picture interactions? Maybe an example with numbers would be helpful in this intro, just give a taste of what is to come and show what can be accomplished with this method.

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

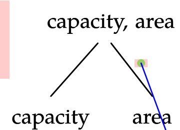

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

So in this example we are referring to leaves as the quarks and electrons? This abstraction is very clear, however, just to reaffirm your point as made in the intro, specifically stating their leaf characteristics would be helpful.

I don't understand what this is serving as an example of. I thought it would be clarified later, but I got lost in the details of creating the tree and don't understand what exactly this example is supposed to represent.

Just curious, what exactly does "local" mean here? Local to the text, or local as in limited-scope?

I think it means local as in it's only looking at a small part of a larger tree? Edit: OK I take that back, I'm not sure anymore.

Local is not the best word here, but it seems to mean specific to the text, as opposed to a grander physics topic like sub-atomic particles.

also confused.

I am almost positive he's just trying to say "an example from this class" or "an example from the text"

Just my two cents, but maybe it would be better to use a different example of abstraction than to describe how to draw the divide-and-conquer trees in the book. On the first read I was confused because we were previously talking about divide-and-conquer as an approximation method and now we're taking it to a new level with abstraction, but the abstraction of how to draw divide-and-conquer trees. Of course on my second read it was very clear, but it did take me a second read.

I understand you want to lay a foundation for not using captive user interfaces, or at least for why they are counterproductive as far as this text is concerned, how ever this paragraph seems very forced and mostly unneeded. You could just cut this entire paragraph and start the next with the first sentence from this paragraph.

Agreed - this paragraph seems entirely unnecessary.

What does this have to do with abstraction?

While a tree as reference is nice since its the basis for the example, its necessity/relation to abstraction isnt defined so maybe an improvement in that area would make its presence more understandable

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

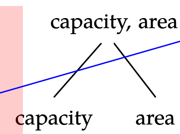

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

this tree bothers me because it is very off center

Should be: tree to the right, as it is not technically in the margin.

Ah, good point, I'm sure most people would figure it out but this should be rectified.

typo. figured if this was to be in your book you should know.

Thanks for finding that. I make so many of that kind of mistake that I once wrote a 19-line Python script to check for double words like 'from from' or 'the the'.

I'll dust it off and run it on the TeX file to see if there are other examples. Ah, there is one more in the next section that you'll get.

I would use a different example.

I agree, this example seems like a odd at first glance and I had to reread it again to understand the application of it.

I think it only makes sense in the context of yesterday's lecture (especially the 'captive user interface' part), and should definitely be explained better or changed.

Do you mean that you have to start over every time you wish to make a new figure? Because you could create a template, even with a GUI program...

Or you could just copy and paste.

it may just be that I kinda enjoy this kind of thing, but I'd much rather do this kind of tedious (but simple) work than thing too hard about how to code it...There is also that whole copy&paste thing is useful for. also peons...peons are nice.

what does captive mean here and why is it a bad thing?

I am also a little unsure of what this reference means.
I believe it means the same thing as GUI, and as discussed in the last lecture, although it is useful, it is extremely limited by what the original designers chose to implement.

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

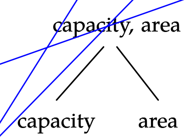

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

You make a valid point about how using Postscript allowed you to save time but I feel that your assertions about GUIs comes across as being unnecessarily partisan.

> I think he is just rying to show other problems through which you can also solve with divide and conquer.

I can see how this would apply to non-extensible GUI applications, but many popular graphical programs now have extensive plugin support. In theory, wouldn't it be possible to write a tree-diagram plugin that would be arguably more straightforward for users with little programming experience?

Here it's good that you compare divide and conquer to abstraction, but the example is difficult to understand. Maybe if this was accompanied by some visual comparison it would be simpler for people that don't understand coding. For example, a divide and conquer tree and its analagy in abstraction.

This advocacy of alternative computational methods seems very one-sided. I agree with later posts that this gets in the way of explaining abstraction. Sussman talks about electrical/electronics components, which I think is a much easier idea to grasp. We can specify an AND gate without having to describe the components every time.

I think I missed whatever link there was between abstraction and tree/coding/captive user interfaces. Why is this relevant? And how is it abstraction?

I am curious about the intended audience for this book. Is it intended for only this class that you teach at MIT, or would you want it to be a text that a similar class at a different university could use? If it is the latter, these programming examples make the book less accessible. I have no programming experience and often get hung up trying to understand how the systems in the examples work rather than relating them to the actual concepts. I agree with what others are saying about actually liking the captive user interfaces because, for my purposes, they work better than anything else, partially because I don't know any programming and would be dreadfully inefficient at any other method.

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

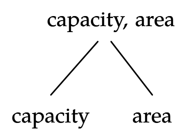

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

Couldn't a text-based graphics abstraction be more difficult for some people to understand? I realize that certain UI's may be limited, and text is one of the easiest forms of input, but depending on the program's audience, couldn't a visual tree-diagram editor be more effective?

I agree that sometimes text-based graphics could be harder to understand. I am personally a visual thinker and graphical analyses always make more sense to me. I have also used PostScript and while the language was easy to use, it was hard to think about and construct a graphic with PostScript. Wasn't it just discussed in lecture the other day how humans in general learn better from visual graphics?

The same arguments have been debated about LaTeX word-processors. While, Microsoft Word uses a more visual way of editing papers, TeX based processors use a more programming oriented way of editing papers. The tradeoff is immediate visual effect to saving time in abstracting to multiple applications or longer documents.

abstraction of a figure- the abstract, divided problem becomes creating one node/edge?

So does that mean that any instance of breaking things down into more manageable subunits is an instance of trees/dividing and conquering?

I don't know much about programming languages, but is Adobe's PostScript cleary the most successful language for this? Shouldn't all programming languages be built from basic fundamental building blocks that are versatile?

no period or capitalize Because

Maybe in the future you could assig a homework that actually uses some programming in either unix or postcript. I believe every MIT student should know how to jot down a few lines of code.

Can we see examples of different diagrams drawn by these programs? Maybe that is unnecessary.

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

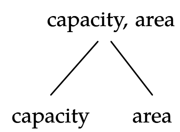

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

what does this mean? is python a full programming language? how is it different from non-full programming language?

I think a full programming language means it has the constructs to make it easy for you to solve a wide variety of programming problems. Python is full programming, but i believe SQL is not

I feel like it might be a better idea to reference a more common programming language. PostScript makes sense, but it seems most people are unfamiliar with it. What about referencing a more common language like Java or C++? One could make a very convincing argument about the use of libraries or packages as a basis for abstraction.

Although PostScript might not be as well known, I feel like it is used as a really good example in this section. In fact, it gives an example of abstraction in a programming language that is tied back to the tree diagrams that we've seen in the previous sections of the reading.

The example is too technical for me to follow. Something more general would be better.

Key word that gets to your point, meaning it can be repetitive, but it is a little misleading, just b/c I have the pieces that are applicable for many complex situations it is not "automated."

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

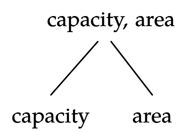

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

This "lazier approach" is still over my head since I am not familiar with programming concepts. I think (just like the last section) that some readers are going to get lost and either get frustrated or just totally skip over the section if they don't understand the programming language.

Agreed. For me, it would be much more work and time intensive to learn how to program to get to the stage where I could do this than just to use a fairly set UI. Especially where there are so many out there that specialized in certain things.

I disagree. Without knowing anything about the languages, I was able to piece together that the alternative was a manual means of making the trees that were copied and pasted. I also understood that he took a shortcut by using a program that had some sort of abilities built in to make the trees.

I think some people have the sour taste of the UNIX example in their mouths, but this is fairly straightforward if you just re-read it a few (3) times.

It seems hard to revert thinking back to 'programming' thinking after having been spoiled with GUIs for so long. I agree that some might skip it because of not understanding programming, but I also agree with the point that it is very important to have that thinking, and that it's a tragedy that we are so keep on using GUIs. Since learning programming, my thinking has become cleaner, and not just for computer skills. It's an improved way of going about life tasks.

Can you post links to where we can learn more about the different coding languages/programs/etc mentioned? I'm not familiar with a lot of them, or have a passing familiarity, and would like to learn more about them.

Good idea. I've added several to the course website. I'll upload it from my laptop to the MIT webserver shortly.

Are we going to continue with a lot of computer programming material? I don't really understand this stuff.

cut mainly...unless you are going to give us a second reason, the word is useless.

I don't have a problem with coding as an example as much as I do biology. I have an unreasonable distaste for that science.

> The importance of this decomposition strategy is not simply that one is dividing the program into parts. After all, we could take any large program and divide it into parts – the first ten lines, the next ten lines, the next ten lines, and so on. Rather, it is crucial that each procedure accomplishes an identifiable task that can be used as a module in defining other procedures.

What they write about programs applies equally well to understanding other systems. As an example, consider the idea of a fluid. At the bottom of the abstraction tower are the actors of fundamental physics: quarks and electrons. Quarks combine to build protons and neutrons. Protons, neutrons, and electrons combine to build atoms. Atoms combine to build molecules. And large collections of molecules act – under some conditions – like a fluid. The idea of a fluid is a new unit of thought that helps understand diverse phenomena, without our having to calculate or even to know how quarks and electrons interact to produce fluid behavior.

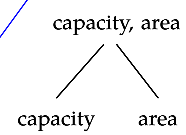

As a local example, here is how I draw the divide-and-conquer trees found throughout this book. The tree in the margin, repeated from from Section 1.3, could have been drawn using one of many standard figure-drawing programs with a graphical user interface (GUI). Making the drawing would then require using the GUI to place all the leaves at the right height and horizontal position, connect each leaf to its parent with a line of the correct width, select the correct font, and so on. The next tree drawing would be another, seemingly separate problem of using the GUI. The graphical and captive user interface makes it impossible to organize and tame the complexity of making tree diagrams.

An alternative that avoids the captive user interface is to draw the figures in a text-based graphics language, for then any editor can be used to write the program, and common motifs can be copied and pasted to make new programs that make new trees. The most successful such language is Adobe's PostScript. PostScript statements are mostly of the form, "Draw a curve connecting these points." because PostScript is a full programming language, by clustering repeated drawing operations into reusable units, one can create procedures that help automate tree drawing.

Instead of using PostScript directly, I took a lazier approach by using the high-level graphics language MetaPost mainly because this language has been used to write an even higher-level language for making and connecting boxes. In the boxes language, the tree program is as follows:

though, I must say that dealing with many languages/programs I haven't yet encountered does add a level of complexity to the homework and reading.

I feel like maybe using energy or something more broad as an intro to abstraction might be useful. I know as a non course 6 student abstraction and recursion were hard for me to grasp at first and a more global example might be beneficial in this section.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

Comments on page 3

This seems pretty unnecessary. Perhaps some pseudocode?

Perhaps you could have little boxes on the sides or before the chapters with code saying what language, which shells to use, and how to run them, in the event that readers want to give it a whirl.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

I have no idea what this means

well, the comments walk you through it. First, the contents of the 3 text boxes are specified, then their locations relative to one another, then it draws the text boxes, and then it draws the branch lines.

This assumes you know that the % before a line means its a comment.

It might be better to write an example (like this one) in pseudocode. While the comments do walk you through it, you can't assume that every reader is going to have the requisite coding experience to realize what exactly is going on.

As mentioned above, I think it would be helpful to "guide" students through the code in this book by explaining in detail what's going on. From reading these posts it seems that many students are confused or discouraged when they see large amounts of code with few comments to help them understand.

While this is a course 6 class, not everyone taking the class has coding experience - why not write a line by line description of what the program does instead of providing code? I would guess that there exist MIT students who have never coded.

I agree that maybe this should be more of a written pseudocode. I understood this because of my programming experience, but someone without may have problems even recognizing what is commented in this code.

It seems like a lot of my thought in reading this is aimed at understanding the code and not the apporoximation method driving the code.

The author can't hold the readers hand at every pond crossing. If the reader doesn't have any idea what the text means (i.e. they have literally never written ANY computer code in their lives), they can still understand it based on the two previous relatively-simple paragraphs.

I agree. While I understand what this code is accomplishing in it's entirety. I don't know what a lot of the specific functions and things mean and I think their presence detracts from my understanding of the concept because I feel like I should try to understand the code.

I am on the same boat about how even though I understand what this code is doing, it is still a new language to me and detracts from my understanding as I try to figure out and picture what each line is doing.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

I have to side with the not explaining it is much,If it's written in psuedocode, it loses a good chunk of the purpose of the code, and it's not intended to necessarily be understood for the coding factor, but for a general overview of the format, supplemented by the comments.

having limited coding experience, i was able to pick this apart but i can imagine how hard this would be for a person who hasn't coded and how important pseudocode would be for the understanding of this. i think that the code should stay but perhaps on the right hand side (as if the upper part of this page is split into two columns) there could be italicized psuedocode or an image showing what's happened after that chunk of code

What about making the code block more of a figure than a part of the text? This way a small caption could be applied with a few short sentences acting as a way to guide the reader through the code, much like psuedocode would do.

I would have just draw it in paint

you don't have to understand exactly waht it meants; only the gist of it. By placing exact code verbatim in the text, Sanjoy makes the application of abstraction seem more realistic. I don't think he should withhold realism from his examples, but rather suggest sources of more detailed explanations for students that need it. In our case, these comment boxes provide a good place to suggest more material and clarifications.

This is both confusing and kind of comes out of nowhere. There has to be a better/easier way to explain abstraction than with this complicated jumble of coding.

Agreed. It would be nice if the code was graphically commented or hashed out a little more. As a Course 2 student, this makes virtually no sense and has compromised my understanding of the rest of this section.

not familiar with this language, i doubt a lot of people are, maybe should have used a simpler diagram creating software, like powerpoint, etc?

It would be nice to have interactive code so we can actually understand the programming better for those of us who are unfamiliar with coding.

Since this will be in a textbook, prseumably, this might be a difficult concept.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

I don't understand several of these lines.

just as a side note, I really do not have any idea how to look at this or any lines of code and get anything out of it, I need someone to explain it to me in person, reading it on a page only makes me more confused

I know that LATEX can be used to make similar PDFs. Would it be possible to do the trees in LATEX (or a higher level tree language based on LATEX) as well?

I understand this paragraph, but I don't see how the coding part above relates to abstraction, or how it is necessary in this article (I'm one of those who doesn't have much coding experience, so I could be missing the point completely...)

The idea is by writing higher level code we can avoid the mess of error prone code that we would need to accomplish simple tasks. By the end of the section this all becomes clear.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

This sounds really confusing. Is this really relevant? Does it help us to understand abstraction or approximation?

I can kind of see how it's an example of abstraction based on how it's an example of defining how different levels of the tree are related to each other.

Agreed... I understand the example but I still don't understand abstraction, and this isn't really helping me understand it at all. It think the examples tend to lose their purpose of being "examples" to demonstrate how a concept works, and kind of become there own thing... which is interested, but not helping me with abstraction. Perhaps giving a short summary of the example, relating it to abstraction, and the delving in deeper?

I think this paragraph is helpful for understanding how higher and lower levels of code work, which is an essential part of abstraction.

I agree; it seems difficult to see how this connects to the "bigger picture"

I think we've seen this same problem earlier in the notes - although some of these explanations are quite interesting in their own right, well written, and quite thorough, they are not always relevant to the discussion at hand.

Yeah it's kind of overwhelming looking at this text trying to understand it and also trying to connect it.

Even though this is slightly overwhelming it does do a bit to show how he is already using abstraction to make tree drawing much easier than using PostScript and slightly easier than using MetaPost by using the boxes language. Later on when he adds a tree language over boxes he attempts to show how adding that level of abstraction makes creating a tree even *more* easy than it was.

While I agree the example as walked through is rather confusing this paragraph does help show abstraction - possibly a better way to represent this would be by some sort of graphic? Possibly a "capabilities funnel" or "lines of code required comparison" of sorts showing the relative capabilities of the methods and how eventually x lines of tree language make y lines of boxes language make z lines of MetaPost and so on

so complicated to explain...

This is a very good observation I had not noticed

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

there's a wrong level of abstraction? isn't the only rule that the smaller sections function as expected?

I am also confused about this. I understand that something can be abstracted to where it is not as useful but how useful or modular does something have to be to be on the correct level of absraction? Why is this at the wrong level?

i think by "wrong" he means that the level of abstraction is not optimal for this example's purpose. i'd imagine that the correct level of abstraction would allow a person to perform their task without an unnecessary amount of repeated text/info. it's probably less definite and more "what feels right", if it seems that there's an unnecessary amount of code for a simple job, it might be at the wrong level of abstraction

Highlighting these repetitions would be helpful for the elementary programmer.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

What? What's a repeating motif? And how does it imply that its at the wrong level of abstraction?

I think a repeating motif is a section of code that is used over and over again with few differences. This "motif" can easily be written as a function in a higher level programming language.

Recursion is another way to abstract away repetition, and it's particularly helpful with tree structures since the same things are happening over and over at different levels of the tree.

I'm not sure because I'm not experienced with coding, but I think what the author means is that if you have code that comes up over and over again, you can write another program or method that handles that code–like in the last reading, the author was saying make each program do one task.

I think it means something that one should use a "for" loop for, in higher levels of program languages. It repeats, so you shouldn't have to copy-paste it again and again

It doesn't even need to be in a "for" loop... sometimes you call the same function at various points in a program, and it is useful to have it as a separate function to call as opposed to repeating those lines of code over and over again... but there doesn't need to be a pattern in when they are called as using a "for" loop implies

I think the previous comment is correct - this doesn't necessarily mean use a 'for' loop, but you should write a program that takes in certain inputs. You can alter those inputs every time you call the program so that the program does what you want it to.

I would argue that a for loop isn't really abstraction and I don't see how a for loop would inherently change the level of abstraction of its specific functionality. I thought of that being done by making a parent class that has common functions for a type of objects. Then the children inherit that functionality without having to rewrite any code. Or something.

I am having trouble making a connection between what I am reading now and abstraction.

I am having trouble making a connection between what I am reading now and abstraction.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

This sounds somewhat awkward to me, or it might just be an idiom that I haven't heard much.

I agree. I thought to "take ones own medicine" means to do something that you told someone else to do, and the idiom isn't very clear in this context.

This makes perfect sense to me, based on context clues. If anything, you might want to change "medicine" to "advice" for people that are struggling.

I understand what you mean, but I think "taking one's own medicine" has a negative connotation from its usage in other places.

Perhaps skip almost everything between where the problem is introduced and the part where you turn graphical trees into outline-type description. Then you can explain how the outline is created by abstracting from the tree. (Though perhaps a better example all together would be better?)

How do these three lines specify the positions and number of the boxes for tree? I don't understand - It says that the preceding tree is specified by only these three lines, I must be missing something.

The top and least indented line goes up top. The next two lines are indented, so that means they go below "capacity, area". Since "capacity" is above "area" and they both have the same indentation, then they go in the same level, but "capacity" goes to the left of "area". In summary, indentation determines what goes below it predecessors, and the order at which you write the words (with same indentation) just go right after each other from left to right. I hope this helped.

this tree-language interpreter was how long in itself? i.e.- was the time spent creating it cost-effective?

It is about 140 lines of Python. It has been very time effective! I just checked my tree directory (if the projector works, I'll show it in class) and it has 34 tree files, one per tree. They total 241 lines. They get turned into 34 files using the boxes language, for a total of 803 lines – and those lines are hard to write correctly.

Furthermore, if I decide to change how the trees are displayed – for example, sideways instead of vertically – I have to change just one program (the translator) rather than 34 individual files.

```
% specify the texts
boxit.root(btex capacity, area etex);
boxit.capacity(btex capacity etex);
boxit.area(btex area etex);
% specify their relative positions
ypart(capacity.n-area.n) = 0;
xpart(area.w-capacity.e) = 10pt;
root.s - 0.5[capacity.ne,area.nw] = (0,20pt);
% place (draw) the texts without borders
drawunboxed(root, capacity, area);
% connect root with its two children
draw root.s shifted (-5pt,0) -- capacity.n;
draw root.s -- area.n;
```

The boxes program translates this program into the MetaPost language. The MetaPost program translates this program into PostScript (or into another page-description language such as PDF). A PostScript interpreter in the printer or in the on-screen viewer translates the PostScript into black and white dots on a piece of paper or into pixels on a computer screen.

Even with MetaPost, a long program is required to make such a simple diagram. A clue to simplifying the process is to notice that it repeats many operations. For example, the direct children of the root have the same vertical position; if there were grandchildren, all of them would have the same vertical position, different from the position of the children. Such repeating motifs suggest that the program is written at the wrong level of abstraction.

After using the boxes package to create several complicated tree diagrams, I took my own medicine and created a language for drawing tree diagrams. In this language, the preceding tree is specified by only three lines:

```
capacity, area
  capacity
  area
```

The tree-language interpreter, which I wrote for the occasion, translates those three lines into the boxes language. The abstraction tower is therefore as follows; (1) the tree language , (2) the boxes language, (3) the

This sounds like an important term, like a tree.. What is it?

Was this term ever defined explicitly as different from abstraction? Is it just abstractions stacked upon abstractions?

I assume that the pixels/specks are the base of our new tower? A diagram, rather than parentheses and numbers, would do well here.

I think it's ok the way it is. I don't think you get any appreciable benefit from a diagram here–the simple list is fine.

I disagree with the previous comment - a picture would be nice, if only to see what's at the "top" and the "bottom" of the tower - visualization and images are more helpful for me.

I agree that if the term is used, there should be a diagram of the abstraction tower both when it's introduced and in this specific example. This would illustrate the hierarchy. The earlier comment is correct that the tree language is the highest level of abstraction, so it would be at the top and pixels at the bottom. Rather than a tower, it could also be represented by a pyramid, because as indicated, the short tree program abstracts away many more lines of PS code.

This description of the abstraction tower is very helpful and this numbered list is an excellent way to summarize the levels of abstraction described in the previous paragraphs. A diagram may be helpful if the reader is not really experienced in the concept of abstraction applied to programming, but I don't find it necessary.

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

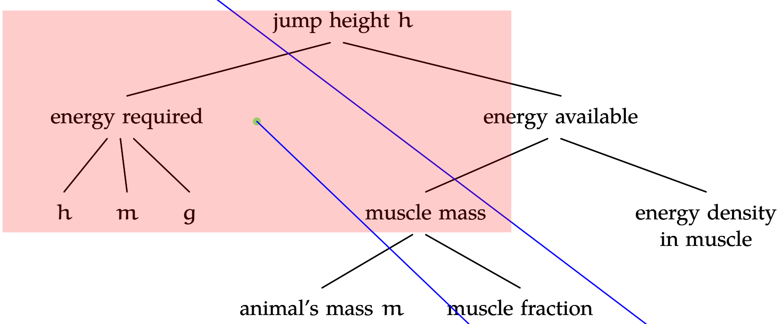

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the boxes language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

The method is a bit clearer here and comes together as you describe how you completed the problem. I think the use of a computer engineering problem is good because it shows the versatility of this method. However, it would help if many of the more technical processes were also summed up in a more simple manner.

Don't understand the last two as well- how can you relate pixels on a screen as being a lower abstraction level than a computer language?

I forgot that we weren't still talking about trees since this example of abstraction is all about the trees that were used to describe divide and conquer.

The coding example somewhat makes sense, but I don't think it should be related to trees since we've been associating them with divide-and-conquer.

Me too. This section has a lot of different concepts in it between the examples and relating to previous chapters. It's hard to keep track of which one we're focusing on.

I agree. Seeing the tree makes my mind jump back to divide and conquer and made it confusing as to what exactly we were looking at.

Why weren't they included in Chapter 1? I would include them and then in Chapter 2 reveal just how you made them.

Ok, I only followed all the following by reading the bottom paragraph. Perhaps rearranging the graphics would help and I lost where our 5 levels of abstraction went...

I couldn't find this example in any of the previous sections. Is this a section that we haven't covered yet? In that case why is it being used without any introduction? This slightly confused me, as I don't remember doing anything involving an animal jumping.

i think this is a later section in the book. if we had the whole thing, i'm sure we could flip forward and check it out.

I agree that abstraction is a powerful tool that can be very useful. However, I think for its introduction, another example would be better to use. Since we just finished discussing trees with divide and conquer, the use of a tree here sort of confuses and blends the two topics together. Understandably the topics are very closely related, but perhaps a different opening example would help differentiate them.

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

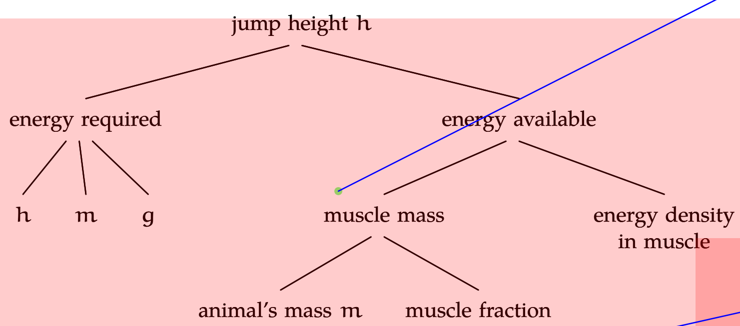

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

I'm not sure what this example is doing here...it's great to see a tree, but we've been seeing them for the whole previous chapter and we know that you can draw them. Without any introduction or qualification, it seems like this example just comes out of nowhere.

I think there might be a better example to start with, perhaps one with less code (although I understand most of the code) and more of a direct comparison to divide and conquer, and how this can be sued to solve problem that more people understand (for example, the oil barrel problem or something similar)

I agree. I think too much coding can easily confuse some audience like me who are not so familiar with coding. I think it's better to start with non-coding problems so that it's suitable for the general audience

And look like python, indentation-wise. It would be nice to mention the indents as an organizational tool!

it bothers me that you are using a tree from a future section to explain this...i'd use the tree from part 1.3 or 1.4

I like the comparison here. It's good to see the tree and also the same in the tree minilanguage.

Why does this h have to be between dollar signs? It's not in a separate tree..

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

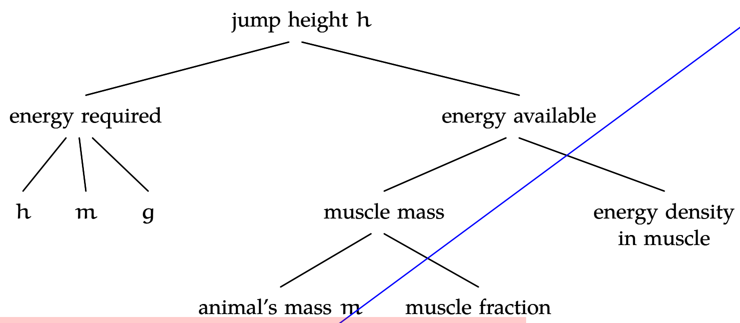

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

what do these $ signs mean? maybe it's better to annotate them on the side?

In TeX, one typesets mathematics by enclosing it in dollar signs (that tells the TeX translator, "Here be mathematics. Turn it into nicely typeset text."). So I adopted that convention here. Thus, h means the variable "h" (in italics). Similarly, $a=bc$ would typeset the equation a=bc.

In fact, the lines of the tree file get passed to TeX for typesetting, so I didn't even have to do anything special to use the math typesetting features of TeX (TeX is a useful abstraction!).

In TeX, one typesets mathematics by enclosing it in dollar signs (that tells the TeX translator, "Here be mathematics. Turn it into nicely typeset text."). So I adopted that convention here. Thus, h means the variable "h" (in italics). Similarly, $a=bc$ would typeset the equation a=bc.

In fact, the lines of the tree file get passed to TeX for typesetting, so I didn't even have to do anything special to use the math typesetting features of TeX (TeX is a useful abstraction!).

I generally found that the examples here made sense, but I didn't feel that it really got to the heart of what abstraction is. We saw a way it was used to create a new programming language, but I think that without any coding experience, the abstraction used there would have been entirely lost on me.

is this is coded tree?? thats kinda cool

I would kind of be interested in seeing this example before the first, because it wasn't until here that I actually fully understood the point you were trying to get across. I'm sort of interested in what would happen if you presented this example starting with this layer of abstraction and then exploring what is actually under the hood.

I see intuitively how this language relates to the tree, but obviously that's because this is some language the author made himself just for this purpose. I do not understand at all how he got from the 1732 lines of PostScript code down to these 10 lines.

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

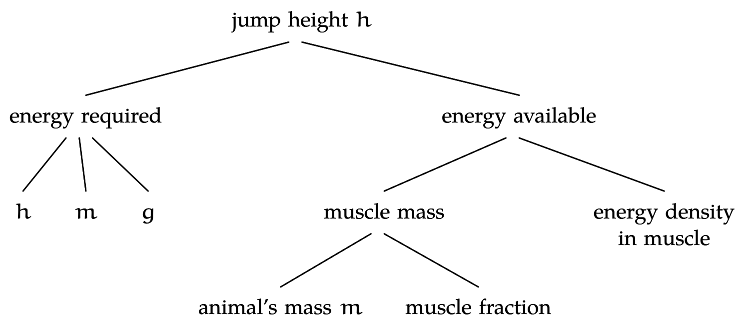

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

I am a little confused as to which language this is in. Obviously it says "tree minilanguage", but I am confused as to whether this is what the author invented, or some combination of all of the other random languages discussed before, a little clarification would be helpfull. In general I would like a little more background when the author claims to invent a new language/method, becuase although it is clear to the author, I often find it confusing, which part he/she invented, why, etc...

I figured the author was talking about the language he invented, but I can see how this could be a bit misleading. However, the important part of this sentence is how the author hints that this process made constructing tree diagrams much easier than before. I don't think the reference takes away from the point that abstraction simplified the problem greatly.

It might not be quite as simple to understand as you think. I think you should explain more thoroughly how the program knows which lines go where on the tree. You should actually mention the indentations in the text somewhere. And you should also specify what the dollar signs indicate.

I think mentioning that you can use the indentations recursively might be useful. Also, you have to notice that the h, m, and g are in a different font (the dollar signs in TeX represent "math mode", and he mentioned TeX earlier... I assume they mean the same thing.)

I agree, if I didn't see the tree right above it I don't think I would have been able to look at this code and draw the tree myself.

what are the dollar signs for? Is there a reason you need these special characters to specify single characters

It's just code, I'd guess it makes the letters italic or in the font that they're in

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

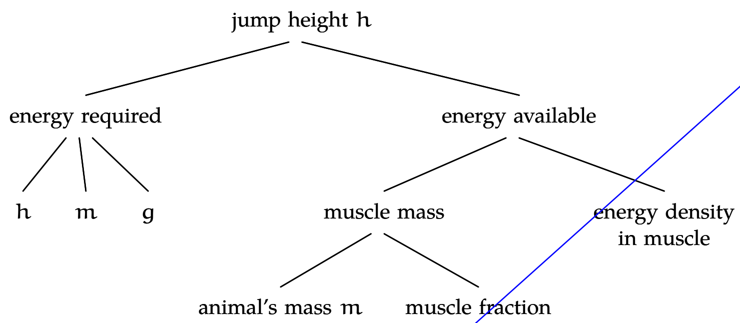

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

I'm probably in the MIT minority, but every time I see computer code I get lost. Don't know what any of it means.

Why, this one simply has the indentations defining the levels of the tree.

I agree with the first comment, if I were just given the code and not the tree, I'm not sure I would have known what I was supposed to do. The tree is helpful and I'm glad it comes before the coding because it makes the code easier to understand. That being said, even code that ends up being more simple like this one is intimidating to look at when you aren't familiar with programming.

Understandably, not everyone is familiar with programming and it can seem confusing to those not familiar with it. However, it is worthwhile to remember that the code presented here isn't in a familiar programming language that many people know, like C, Java, or Python. In the previous page, it is mentioned that this "language" was made up by the author. So, I have never seen this language or code before, and yet I can analyze its content and structure and make guesses about what it means. I feel like a lot of people in this class see code and panic and forget to spend the time trying to figure it out. If you encounter a new mathematical or engineering formula, do you panic because you've never seen it before? No, you slowly go through it term by term until you understand the whole thing.

Uh, that said, coding isn't so much about learning a language so much as a way of thinking... so if you've never coded before this is definitely hard to understand as an example, especially when you're still trying to wrap your head around the idea of abstraction. I guess what I mean that it doesn't help if the example you use to explain a confusing topic is also confusing.

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

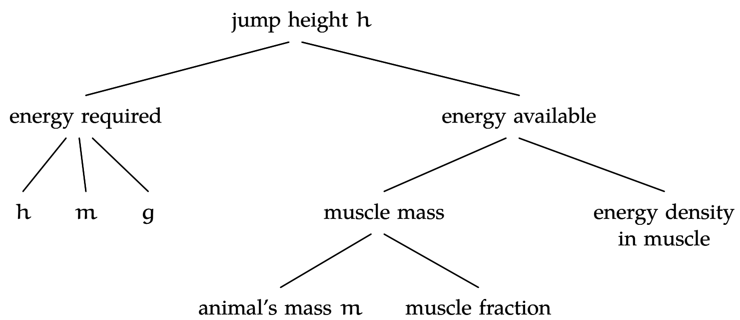

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the boxes language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

I agree with the previous two comments. I think Anon1 is simply pointing out, as I have also observed, that any time there is a snippet of code, there are at least 5 comments about how it's confusing and asking why the code is there. I think this is something perhaps Professor Sanjoy should address directly, but I would like people to remember that this is a designated course 6 and course 2 class, both of which require you to have some exposure to coding (matlab in course 2), so really, there's no reason people should be so surprised by a few lines of some meta code (or even less so of unix command line prompts in the previous section). That said, I do grant Anon2 that reading and understanding code comes much easier to some people than others. If you're having trouble with this bit of code, just think of it as a list. Everyone knows how to use lists to outline ideas, and this is no different. The further indented the line, the more nested and subsidiary the element.

I guess this is part of the reason the class is in Course 6 as well. I think it's good to have Course 6 and Course 2 examples. But I am tempted to argue that Course 2 ish examples make sense to Course 6 while Course 6 examples make less sense to Course 2 students. What do you think?

Overall, the example is pretty easy to follow, but I'm not sure how well it actually illustrates -how- to abstract. It pretty much just says: graphs are hard, so we abstract to this outline form through (black box). As long as a later section tells us how to practically use abstraction, this is fine.

Why does one leaf have a variable and other don't. I feel like each end leaf should consist of a variable. Also are the indents defining whether a new node is made?

This is a much better example than the capacity, area one.

I think it confuses me more actually... maybe since the capacity one was something I could better visualize.

but since this minilanguage is not real, are we merely to take from this that we should be writing as elegant programs as possible (i.e. dividing and conquering with intelligent redundancy?)

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

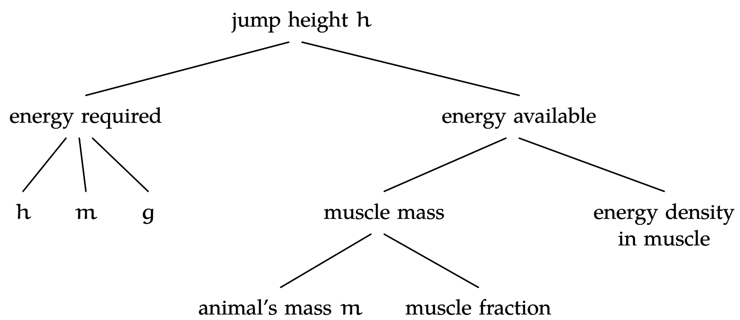

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the boxes language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

why is the boxes language error-prone? Is it always error-prone, or just in this particular case?

I think he was just pointing out that writing those 34 lines of code by hand would most certainly introduce errors the first time through, whereas using a layer of abstraction (his metacode), he was able to create those 34 lines in a fasion that did not produce errors, a sign of a good abstraction.

That is cool.

Agreed, this is a great closing paragraph and helps helps point out the now obvious benefits of abstraction.

Excellent closing though I feel that maybe this sort of growing code cost could have aided earlier in understanding the long list of boxes over MetaPost over PostScript. I do agree that this really nails the benefits of abstraction

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

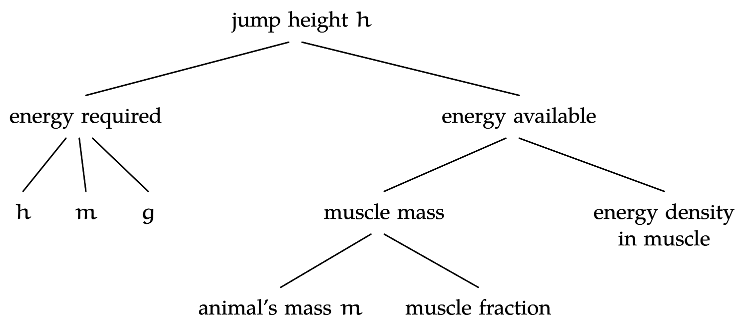

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the boxes language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

That's pretty cool, but I'm still waiting to see how this chapter and the last fit into the scope of this course.

The last chapter discussed divide and conquer and this one is discussing abstraction?

I like the way the divide and conquer section was set up with the concept explained with an estimation example and then used the code as an elaboration on the topic in a different section. I think a simpler "warm up" explanation of abstract would be good before adding the complexities of coding.

Agreed - it might have been better to begin this chapter by framing abstraction as useful for approximation. As it is, I think many people will read this chapter with the impression that abstraction is useful for coding.

I agree. When the chapter began with abstraction, I was thinking approximation methods that can be abstracted to solve different types of problems. All this discussion of GUIs lost me, and I still don't understand what it has to do with abstraction.

I agree. I understand that this is the beginning of a new chapter and a new concept. However, in the introduction, it was confusing about how "abstraction" related to divide and conquer, where it fit into the picture.

Again, I agree. I think this in part stems from the fact that this new chapter begins with a definition of divide and conquer.

I agree that a little explanation of why abstraction is an important technique in approximation could have been useful at the beginning. Other than that, I thought the PostScript to tree minilanguage example was an excellent way to introduce the concept of abstraction especially to people who are not familiar with the concept from programming experience.

i love this explanation...really makes your point well and it's a little scary to think about...but cool

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

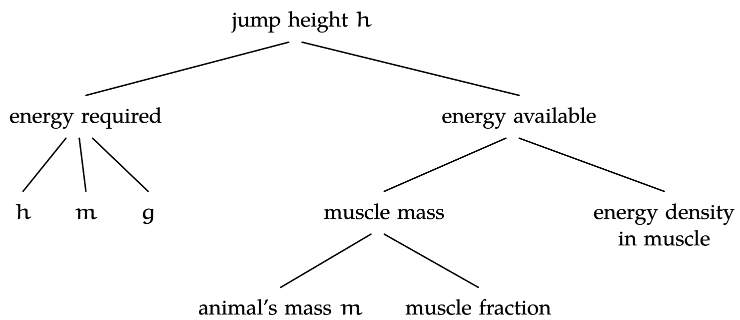

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams

A powerful kind of abstraction is a diagram – for example, the trees illustrating divide-and-conquer reasoning in Section 1.3. Diagrams are

this seems pretty inefficient. Aren't we normally striving to make things simpler and efficient? Also I am a visual person and in order to write a tree in tree language I would have to first draw it out on paper. Does this mean that the tree language wouldn't be very useful for people like me and that using a GUI would be better?

it's efficient if you're trying to visually represent your hand-drawn tree on paper. imagine how else you'd be sending the document to a friend without a scanner at hand. you'd probably use word or ppt to draw boxes and draw lines between them. or you could use this concise program to do the same thing with no drawing, just spaces and words

I didn't understand hardly any of this section. The example gets too caught up in the programming details without referring back to the abstraction concept often enough.

I agree- I don't feel like the idea of "abstractions" was represented- although that might be my own confusion

Although I can see this example as a demonstration of abstraction, I agree and think there was a little too much coding and specific references to unfamiliar languages... at least for an introduction. While it still greatly applies the principles of abstraction, it might be better suited for a subsequent chapter, where we get into a little more detail explaining abstraction.

The main idea is that programs one can write a program to take an input and do something with it independently of your main objective. Hence your solving two seperate problems.

this sentence is awkward...use 'diagrams are a powerful form of abstraction.' ?

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

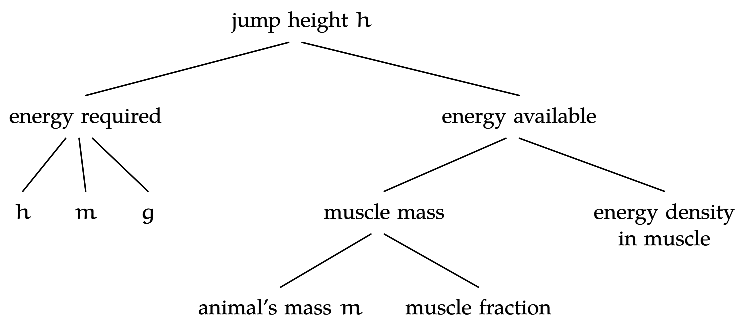

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the boxes language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams as abstractions

Diagrams are themselves a powerful kind of abstraction. Diagrams are an abstraction because they force one to discard irrelevant details, reducing

This explains abstraction a lot better. Before it seemed like you were creating things to make life easier, which doesn't help estimation from my point of view. Now it is clear that you have to think about problems differently sometimes.

I am having trouble with replying in the notes section. When I hit "notes" and want to reply, I go to the options and the drop down menu comes out nicely. Then, when I hit reply, the box comes out just as expected. However, when I try to click anywhere inside the box (in case I want to fix a typo or anything) the same drop down menu comes up again. I just realized a lot of my previous nb posts were never posted because I didn't realize what as going on...

Also, it happens within this box as well.

This example is a lot easier to follow and I think it illustrates the concept of an abstraction more easily than the UNIX example. Maybe if you gave this section first and followed it with the UNIX example it would work better.

It seems that abstraction would be much harder for a problem in three dimensions.

It seems to me that this version of abstraction is easier to understand than the previous example given in Wednesday's reading.

So, is this kind of a correlation between visual diagrams and memorization. I wonder if that helps with divide and conquer or when you have to remember certain constants/figures for abstraction

I am a bit confused as to why this is a downward schedule. What is a downward schedule? Are you thinking of time as cyclical?

seems like a pretty interesting research topic...course 9 perhaps? wonder if other primates have intuitive estimation skills (moving from tree to tree, for example)

is this referring to something like the Rorschach test?

like a previous commenter, I am wondering what you mean by considering the Navier-Strokes equation as symbolic information...are we supposed to look at it as a picture and remember it ie. photographic memory-like?

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

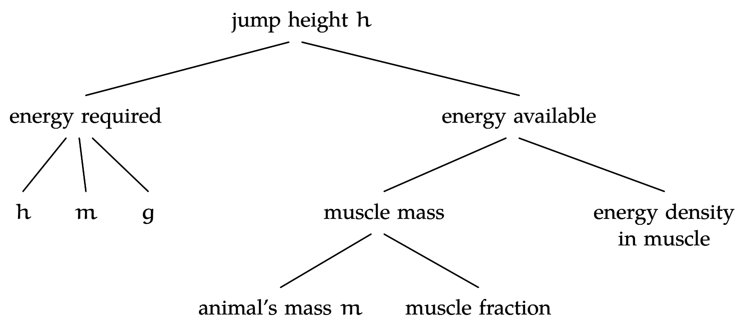

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams as abstractions

Diagrams are themselves a powerful kind of abstraction. Diagrams are an abstraction because they force one to discard irrelevant details, reducing

was the question of the constant slope ever resolved? I agree that this is a confusing diagram, I don't think walking up a mountain and resting along the way corresponds to a straight like as depicted.

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

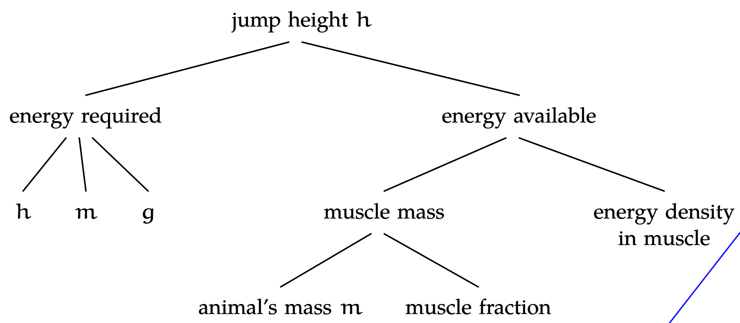

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the boxes language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams as abstractions

Diagrams are themselves a powerful kind of abstraction. Diagrams are an abstraction because they force one to discard irrelevant details, reducing

Read Section 2.1 and do the memo by Thursday at 10pm. Hopefully a change of pace for those weary of programming!

I agree that this was much more accessible than the previous chapter - programming does limit your audience, while everyone can relate to and understand the mountain problem at the end of the chapter.

While I did not find the previous section(s) inaccessible - despite not having tons of programming experience - I found this section to be extremely clean and compact, and good for demonstrating abstraction of data and graphical representation. I don't really think the complaints about the graphs are valid, either, as they serve their purposes without being overly complicated.

To me this is the key to abstraction, look at the big picture and don't stress over the details.

does this mean that tree diagrams are a form of abstraction?

MetaPost language, (4) the PostScript language, and (5) pixels on a screen or specks of toner on a page.

The tree minilanguage made constructing tree diagrams so easy that I created many diagrams to explain divide-and-conquer reasoning in Chapter 1 and to explain the subsequent ideas in this book. Here is a figure from Section 4.4.1:

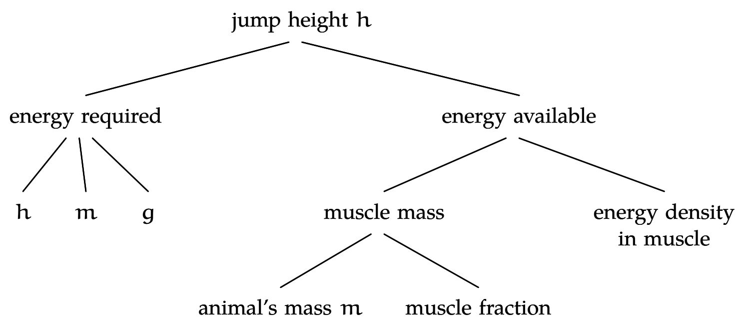

Its program in the tree minilanguage is short:

```
jump height $h$
  energy required
    $h$
    $m$
    $g$
  energy available
    muscle mass
      animal's mass $m$
      muscle fraction
    energy density|in muscle
```

These 10 lines – simple to understand, write, and change – expand into 34 lines of tedious, error-prone code in the `boxes` language. And they expand into 1732 lines of PostScript code! As Bertrand Russell said, "a good notation has a subtlety and suggestiveness which makes it almost seem like a live teacher" (quoted in [23, Chapter 8]).

2.1 Diagrams as abstractions

Diagrams are themselves a powerful kind of abstraction. Diagrams are an abstraction because they force one to discard irrelevant details, reducing

This is how I usually think about abstraction. Today I was a bit confused when abstraction=repeatability. For me its about discarding details and broadening definitions.

I agree–and this is also how you explained this in the very first lecture as a way of organizing your desk. Just throwing some useless stuff away.

Is this the idea behind discarding certain details (for instance, in the MIT budget example, we considered the wages to be by far the most significant contributor)? How should we judge which factors are irrelevant besides general intuition?

I'm also confused as to when we should consider a factor irrelevant. It seems that by using certain abstractions we may be tempted to add even more "irrelevant" details to the problem.

I think there's a difference between neglecting factors and discarding details. Abstraction is the latter. When we talked about wages, it was abstraction when we discarded the detail that people's wages differ. It doesn't seem like abstraction to assume that the budget is made up solely of wages.

This definition doesn't really line up with the re-using definition we have been talking about. Re-using seems like it would be more modulating than abstraction.

I agree. I'm confused why we are considering separating tasks and discarding details as the same procedure.

I don't think abstraction throws useless stuff away. I think abstractions just sums up the unnecessary details into one compact thing. I used this example before but its kind of like a circuit element. Instead of using mu*A*N^2/length we simply use L for inductance. Abstraction would be using the L instead of the big long equation.

That's the same way I think about abstraction. My personal example is writing MATLAB and excel functions formulaically so that the values can be easily changed instead of always having to retype the equation or calculation.

I look at abstraction as a way to take the information we know to solve a larger problem. I try to throw out small details by changing the information I know.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla} p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

Comments on page 2

I feel that this should read "our brain's visual perceptual hardware"

They

I also feel like using diagrams can leave out a lot of necessary details and confuse the readers and create misunderstanding.

too bad. seems like the symbolic-processing is a lot more prevalent at MIT

I really like this description of diagrams- I feel like it really ties back to the example in class where we were much better able to understand powers of 2 through a box (1ˆ2 + a top and side layer= 2ˆ2 + a top and side layer= 3ˆ2 and so on).

This is a lot easier to understand description of abstraction as well. For some reason the ones in the past section didn't work as well for me.

Is this true for everyone, despite what type of learners they are? It strikes me this would vary between auditory/oral/written learners, etc.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla}p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

I feel like this would be more powerful if you supported the point with an example demonstrating it.

It's interesting because we talked about this earlier this week in 6.UAT. We were shown a concept in signal processing first with equations and math, and then graphically. Everyone voted and overwhelmingly agreed that they understood the graphical representation more clearly and preferred it.

We were then asked again to describe the difference of how we felt between the two methods. Overwhelmingly, people described the mathematical explanation as more rigorous, complete, etc.. However, Prof. Freeman then said that in this case, we were wrong, the graphical explanation actually was more rigorous, it described more aspects of the signal than did the math, and from it we could extract more information!

It seems that people attribute rigor to mathematical explanations more easily and give it more credibility, even though people tend to perceive and understand graphically explanations more quickly and easily.

A famous example along the same lines is Feynman diagrams. They are used throughout almost every field of theoretical physics now, but they were first invented by Richard Feynman (MIT Class of 1939) to calculate otherwise almost impossible integrals in quantum electrodynamics (roughly, quantum mechanics plus electromagnetism and special relativity).

The previous methods involved horrendous symbolic calculations; but Feynman made it so easy. One of the fellow winners of the Nobel Prize is reputed to have said, "Feyman made it so that any fool can do quantum electrodynamics."

I think it's more than just human evolution. Within our own lifes, we are taught at a young age to learn with pictures, and we don't learn until much later how to formally express our ideas in words.

This might be a good place to include your example of the power of diagrams in class; with the area of a square equaling the sum of sequential odd numbers.

Interesting.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla}p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let u(t) be the schedule for hiking up the mountain, and d(t) be the schedule for hiking down the mountain. In this representation,

I think showing an example of solving the following problem about the hiker without a diagram would be more effective at proving that one method is easier than this description is. Going into the history of how we think seems kind of random although it is interesting.

I agree that it's a bit of a digression, but it's definitely interesting and sort of frames the argument for why diagrams are powerful abstraction tools.

I thought I was reading my psych or linguistics book for a minute there.

Do you think the human capacity for signal processing will catch up with the human capacity for visual processing?

I don't think so. A lot of signal processing still uses vision and so visual processing would be strengthened even as signal processing is strengthened. This may not be absolutely the case, but to support the distance between the two, there is actually a term in neuroscience - the Pictorial Superiority Effect or PSE. PSE specificlally relates to memory but has the same gist as what is described here in the text. Check out Brain Rules by John Medina

(10ˆ5 vs 100,000. Nice touch...)

I like this section. Examples for why things work better tend to help people understand what's going on and why they should care.

This explanation is a really interesting way to think about why visual learning is more effective....although I'm unsure how valid an argument it is

A perfect example of this is the exmple in class of the sum of odd integers and the square

I can't remember if this was discussed when diagrams were first introduced, but I think it would be more appealing to see it the first time we discuss diagrams, because it is definitely the more intuitive rendering of knowledge for a lot of people.

I am a little confused about how sequential analysis leads to symbol processing. I would have thought that symbol processing is more related to visual processing than it would be to language. A little more explanation about this distinction would be helpful for me.

couldn't you also say then that there were limitations to diagrams that required the use of symbols- yes diagrams make things easier but most difficult problems can't be solved using them?

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla} p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

I love how much random information we learn in this class.

I agree a million times.

use: "(even longer among vertebrates)"

Not exactly sure what you're getting at with the line about "general perceptual processing". Do you mean that because at some point our ancestors could hear a predator rustling in the branches, we're better at solving problems today? Is this a reference to our "gut feelings"? Or is this just telling us we've been thinking critically in some way for a very long time (I'm not sure that's even true...)?

We went over this idea in class a few lectures back. Perceptual refers more or less to visual or organized in a way conducive to that.

I think what he's referring to is like the example we did in class where we have memorized the squares of numbers but we don't necessarily know HOW we know them. Then we looked at that box that started at 1 box and added boxes to the top and side to get the squares (ex. 1 box + a top and side row = 4 boxes (2^2), then plus a top and side row= 9 boxes (3^2), etc.)

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla}p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let u(t) be the schedule for hiking up the mountain, and d(t) be the schedule for hiking down the mountain. In this representation,

It would be fun if this diagram was actually included in the reading with a short explanation afterwards of what the picture is - or are we doing this in class?

Does it have copyright restrictions? That's the only reason I could see for not including it.

http://www.bbc.co.uk/radio4/reith2003/images/lecture3_dog1.gif I think that's the one.

Alas, almost everything has copyright restrictions now. I could probably get permission to include the diagram, if I were publishing the book like normal; but I am planning on publishing it (in print and online) using the Creative Commons NonCommercial ShareAlike license, and I doubt any commercial publisher would allow "their" photograph to be licensed that way.

Since the terrible Supreme Court judgment in Eldred v. Ashcroft, which allowed Congress to extend copyright terms as much as it wants, this problem is only getting worse.

In this particular case, there might be a way out. I know Richard Gregory, and it's possible he still has the copyright to the picture, in which I think he'll be happy to give me the necessary license. I will ask.

Meanwhile, I'll show it in class one day (maybe Friday).

I think there might also be other versions that are not copyrighted? I've seen the image though and it is rather striking that once you've seen it once...you cannot go back to the original blotchy view!

I looked at the link several posts earlier, but I did not see any sort of figure. Is this the right picture? Is there anything special that we should look for?

Where is the picture? I want to see it!

I agree. Without the picture this is just confusing...

"...not to spoil the surprise." Clever technique to peak interests of readers.

I'd like to see this...or you should come up with an example that you _can_ print in your book.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

Don't images / diagrams decay over time? where as language or equations must be more enduring to be structurally sound?

I think his point is that you will remember the hidden picture if you see it again. He isn't referring to the physical decay of the picture over time.

also, in terms of equations, if you haven't used it for a long time, you will tend to forget about it.

Yeah, I think he is talking about the first impressions and the impact on the brain.

Perhaps it would be better if you used a different visual then, one that is more well-known.

ooh that makes sense

This paragraph may be a little unnecessary. Going directly from the previous to the next may prove to be more fluid and natural. Also, there will be no chance for people to complain about not understanding the Navier-Stokes equation.

what is the connection here? how does the picture connect to Navier-Stokes equation? the transition isn't so clear to me

The arguement could me made that the primary reason that this is harder to understand is that the meaning of these symbols is defined by someone else. If we had a deep understanding of the symbols, perhaps the equation would be understood almost as quickly

I think presenting the equation is distracting. It draws the eye, placing too much emphasis on a minor detail for a broader explanation. For your explanation, it's not important that people know what Navier-Stokes is. Therefore I think you should omit the actual equation.

well, no, I sort of disagree with this. It would've been nice to have a picture alongside the equation just to point out that it is easy to remember photos but difficult to remember equations, but it reminded me of how I often have trouble remembering equations while I usually remember photos I've seen. Then again, I guess I don't have to recall them as often.

This seems a little random, what does it have to do with the topic?

what is the difference in learning symbols and pictures? do they use different parts of the brain?

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

Even with math (say linear algebra), it is easier to deal with R2 space because you can simply draw things out easily as opposed to R3 or RN

jumpy. either cut this line or make it flow from the previous & into the next sentence better.

Maybe have the equivalent diagram that would make it easier to understand, to compare the two medias?

Like another reader, I've never seen this equation before. Can someone give a brief explanation?

According to Wikipedia, it is a fluid dynamics model type equation. I too would have liked some sort of BRIEF phrase describing it (or perhaps some other more easily conceivable example), though I'm against removing it. The point is to illustrate some confusing concept that required many presentations, and that is one such equation.

As mentioned, its the fluid flow equation for things that have pressure and velocity (like your arteries). I also would vote to keep it, although I know the equation its also a nice demo of many symbols which, despite having been seen before (or never) nonetheless look confusing and easily forgotten

It is used in fluid dynamics for describing fluid flow. It takes into account the various forces which are at work and allows you to full characterize (describe) the flow of the fluid. However, it is a fairly complex equation and often can only be solved analytically in a a few cases.

It does not matter what the equation is - it's the idea that it has a lot of symbols that are harder to process than pictures. The fact that you don't get it probably just proves his point.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla} p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

Would it be more useful to use an example that more people in the class would be familiar with? I for one havec no idea what this equation represents, so I can't really relate to the example.

Now that I've read the whole chapter, this example actually seems entirely unnecessary - I would actually vote for scrapping it entirely.

Agreed that it seems unnecessary. Said by someone who has seen this equation a lot. And for me, learning it didn't require many presentations but many uses (ie I didn't learn it from reading and lecture but from the problem sets and examples)

The equation just serves as a counter example to the diagram argument. The more confusing the better in my opinion. The point is to convey the complexity of symbolic information.

I think he should at least say what this equation is used for. I've never seen it before.

Navier-Stokes equation is used to describe fluid flow.The wikipedia page has a pretty good description:http://en.wikipedia.org/wiki/Navier-stokes

I think he was just trying to say that, generally speaking, it can be harder for us to memorize symbolic information like formulas as opposed to visual information. But I do agree that, by not knowing this formula, I lost some of the meaning of this example.

Just mentioning that it's "the Navier-Stokes equation to describe fluid flow" in the text would seem to alleviate this problem.

There have been a lot of examples that were a bit harder to get through and were completely unfamiliar, but I think that's one of the appeals of this class. These are things that I would probably never learn in my course, but a small does of it is fairly interesting.

The author isn't teaching this equation to us! The point is to show the reader an example of how visual perception differs from symbolic perception. All in all, I think the examples definitely prove the point, no matter what the readers background is.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla}p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

I have to say as someone course 2 I can't help feeling a bit of glee: "Ha, after all this UNIX and course 6 material, finally something we course 2 people can understand and those course 6 people can't! Hahaha!" That said I'm not sure learning an equation is a visual thing for most people, so maybe this isn't the best example.

I think the author's point in including this seemingly obscure equation was to show that our mental hardware is more devoted to looking at diagrams rather than symbolic information–in this paragraph he made the comparison of Richard Gregory's splotch picture to this Navier-Stokes equation

you must be course 6. the course 2 people had no freaking clue what the unix stuff was about.

lol, so true!

i think there is even more mental hardware devoted to cognitive reasoning than visual processing.

Is this referring to the parts of the brain like V1 and somatosensory processing areas?

Not sure what the author means either, but is it necessary to include this? How about just starting from "A good problem-solving ..." I think the author already makes a valid point in the previous example about how the brain's visual perception is very unique and powerful.

probably but more generally just that we have a lot of infrastructure in our brain set up for visual processing

I'm wondering the same thing. What is this referring to; can we get a note? Not necessarily in the reading, but just for our knowledge. I'd like to know if this is referring to a specific part of the brain, like V1-V4, or the prefrontal cortex, etc.

So I just think that making diagrams are a waste of time. Maybe they are useful for teaching others that don't understand the question, but I don't see how making one yourself will let you learn better.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

I could imagine that this is not universally true. Are there some problems where I diagram might mislead you or not be the optimal problem solving strategy?

Although a diagram may not be the most efficient way to solve every problem I think by definition a good diagram should not mislead or confuse you

There probably are, but the author just says here that "a good problem-solving strategy" is to use diagrams. The purpose is just to make things simpler, since in a diagram you only include necessary info. So if in a problem you had a simple, easy to see situation that needed lots of calculation, a diagram wouldn't help.

I think most of us already do this, although I sometimes make the diagram in my head.

I agree, we did this a lot in 8.01/8.02

Its like block diagram representation in any signal processing or control class like 6.003 or 2.004. We don't necessarily know the implementation of the blocks but its the first step to figure out the solution.

While a lot of us already do this I think it's beneficial to explicitly state it since some may not realize they are doing this. Additionally for anyone that may also be interested in teaching of any kind, stating this is a good reminder that visual learning is pretty powerful.

While a lot of us already do this I think it's beneficial to explicitly state it since some may not realize they are doing this. Additionally for anyone that may also be interested in teaching of any kind, stating this is a good reminder that visual learning is pretty powerful.

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do the problem solving using the diagram – in other words, make an abstraction and then think of the problem in terms of it. As an example, try the following problem:

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

is a tree an abstraction then? what other types of diagrams are we talking about?

I think the whole point of this section is that a tree IS an abstraction...a visual one

But it's also important that there are more possible visual abstractions than trees, such as graphs, charts, etc. that can help one understand a problem or result.

So any pictorial diagram can work, as long as it simplifies the problem (enough to generalize a situation without losing "key" details)?

I would say a tree is a good abstraction. If you've ever taken 6.041 or an equivalent, the probability trees are incredibly helpful when trying to look at scenarios.

Not sure whether it technically is or isn't. I would say that it's important for it to be a reusable unit for efficiency in our approach, like the number of poeple in the US.

I also thought this was intuitively yes, and then realized I had no good explanation why. The graph at the end was very nice in explaining this.

This is a great problem - it's appeared on interviews multiple times for me.

I agree. I remember I first encountered this problem reading Marilyn vos Savant and found the solution fascinating. I've also found that several companies use it as well in interviews.

I also found this problem and the solution below interesting, but I felt like I wasn't that interested in it because it was merely a mental exercise. Would it be possible to have an example that had more of a real world application?

In regards to the real world application, I think the fact that this question appears in interviews serves as an indicator of its usefulness. It tells the interviewer how you think, whether you can think this way. Personally, this class is teaching me different ways to think about problems and this was one of the many interesting examples we've encountered.

i really like this problem also. is there any way we could have numbers associated with it at the end to give us something more to imagine (the different paces, etc.)?

I think it's interesting how on interviews a lot of companies don't let you write/draw diagrams in these interviews. It's like, anybody can figure it out if you draw it, but they want to know who doesn't need the visual, which seems a little odd to me.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla}p + \nu\boldsymbol{\nabla}^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

Why is it mentioned that you sleep for 24 hours at the top? Wouldn't it be the same if you just hiked back after reaching the top in 24 hours?

Since it seems a lot of people have already heard this example or think it is intuitive, it might serve well to have a different problem that is more difficult or has more distractions that could confuse people for the diagrams- I would say in this example it's clear that we only care about the distance on the path and the time so there aren't really any distractions in the problem. In other problems, sometimes the hardest part can be determining the relevant information so a question that requires that would be a good fit here.

I disagree. I think this is a great example and works really well here

I also agree; if anything, it reminds me of the word problems we used to do in elementary school math that I had a lot of difficulty with, and is great evidence for me that trying to process this via language rather than visually is much more difficult.

I'm trying to edit this to "I disagree" but I seem to be having trouble editting...

I disagree, I think this example is fine and even if we know the answer, seeing how what we thought in our heads translated into graphs was useful. at least here the problem doesn't detract from the concept

I also disagree – I've certainly heard this question and its solution before, but I think it's much easier to think about in a diagram format.

wouldn't you just rest as different times this way you would be on a different schedule.

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\mathbf{v} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

I found this to be a pretty quick problem to solve intuitively based on a simple continuity argument (which is, I guess, demonstrated by the diagrams). I feel like it would be more helpful to either in place of this example, or after, give a problem that can't be rationalized in our head because it contains several possible interdependent situations, or some other complication.

I agree.

Well, not necessarily. Examples are meant to be super simple cases, the concepts of which can be applied to harder and more realistic problems. Here, I'd rather have a nice and easy to understand example with basic diagrams to prove this point rather than one that will lose me along the way (like UNIX code!). Plus, we're MIT students, so we don't count.

Haha, I think we count, but it's nice to have something that makes sense and can be explained to demonstrate what we're studying

I agree, the simpler problem let me focus on what we were actually trying to learn rather than trying to understand the example so i could understand the concept and getting lost.

It is also nice to have all the information that is needed in front of you so you don't need to try to delete information in your mind.

I like this! Often times questions contain useless words/data that aren't needed and usually end up confusing me

I agree. This problem was so much easier using abstraction. I would like to see another example of this in class.

I take back my last comment. This is a good example of when a diagram is useful.

wouldn't it make more sense to thing about the details you do care about and then ask yourself if you really do care about them?

When I first thought about this problem, a visual did come to mind, but it was not a graph, like you presented. It was a very literal mental image of me walking up the mountain.

since it's a 24 hour schedule. Don't you have to pass where yourself at some point?

a problem to what can be taken in at a glance. Diagrams are powerful because our brain's perceptual hardware is much more powerful than its symbolic-processing hardware. There are evolutionary reasons for this difference. Our capacity for sequential analysis and therefore symbol processing took off with the advent of language – perhaps 10^5 years ago. In contrast, visual processing has developed for millions of years among primates alone (and even longer among vertebrates generally), and general perceptual processing is even older.

Because of the extra development, visual learning can be rapid and long-lasting. For example, once you see the figure in Richard Gregory's famous black-and-white splotch picture [22], you will see it again very easily even ten years later. (I am being obscure about what the figure shows, in order not to spoil the surprise.) If only we could learn symbolic information as quickly. Although I now know the Navier–Stokes equation,

$$\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\boldsymbol{\nabla})\mathbf{v} = -\frac{1}{\rho}\boldsymbol{\nabla} p + \nu \boldsymbol{\nabla}^2 \mathbf{v},$$

learning it required many presentations!

Given the massive amount of mental hardware devoted to visual processing, a good problem-solving strategy is to translate problems into diagrams and do a lot of the problem solving on the diagram – in other words, to make an abstraction and then to think using it. As an example, try the following problem.

> You hike a path up a mountain over a 24-hour period, resting along the way as you need. You sleep for 24 hours at the top. Then you walk down the same path over the next 24 hours. Were you at any point on the path at the same time of day on the way up and the way down? Alternatively, is it possible to walk up and down on a careful schedule to ensure that there is no such point?

It is hard to solve without making a diagram. To make the diagram, first decide on details that you do not care about – consider the situation "independently of its attributes." For example, the day of the month, the year, or the age of the hiker are irrelevant to the solution. All that matters is the schedule on which you walk up and down, namely where are you when? A particular walking and resting schedule can be abstracted to a function of t, the time of day, where the function gives your distance along the path. Let $u(t)$ be the schedule for hiking up the mountain, and $d(t)$ be the schedule for hiking down the mountain. In this representation,

this is assuming a constant speed taken throughout the climb yes? if yes, what would happen if you were to vary the speed going up and down the mountain and how would the graphical representation look in this instance?

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

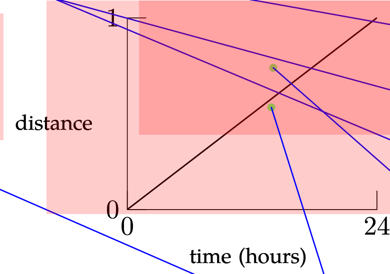

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

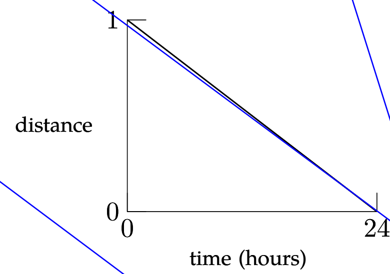

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

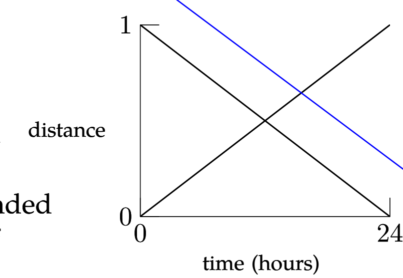

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

it would be nice to have figure numbers so that you can reference the figures easier in the text. (this could be my own bias from reading papers a lot recently)

I think you should specify what you mean by t here. It's not absolute time, since the events happen on different days; rather, it's time since the start of that particular day.

Are formulas also abstractions?

I like how you turned the problem question into a mathematical concept. It helps to understand what we are trying to do witht the diagrams

I suppose if you really wanted to convince people you could show multiple graphs, like exponential movement or something else that wasn't simply linear. I know the result is the same, but it sometimes helps to see multiple graphical displays!

So the diagram represents the abstraction in this case, or is it the abstraction?

At the risk of being even more convoluted... I think it's an abstraction of an abstraction.

So, to answer your question. Both?

It does seem like we're dealing with two types of abstractions: a problem solving abstraction (the diagram) and the data abstractions themselves. So I agree, sounds like both are correct.

i didn't know that was a word.

It is a word, but this seems less like a diagram (a simplified drawing of something like a tree or a schematic) and more like a plot. I think "graphical representation" is perhaps a better term. It's very similar to finding the intersection of two linear equations graphically rather than solving for it. We wouldn't call that diagramming.

I felt kinda dopey when I realized that the answer to people's requests that the graphs should be non-linear was so simple: the linearity of the path is an *abstraction* of a pace that is probably full of curves. We abstract so we can simplify and solve.

I think these drawings are missing the flat parts that you mentioned in the reading

Yeah I agree. the description describes a very non linear walking schedule. First walking fast indicates a larger slope than the walk after the rest. and the resting part is missing.

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

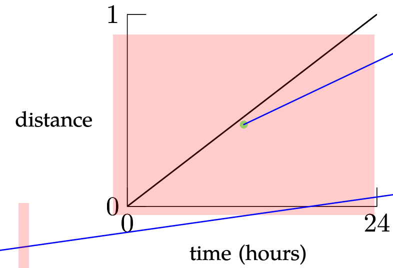

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

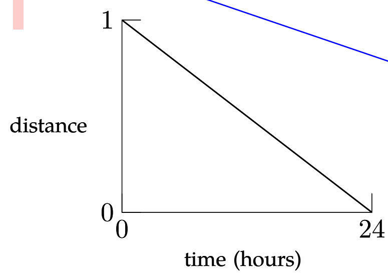

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

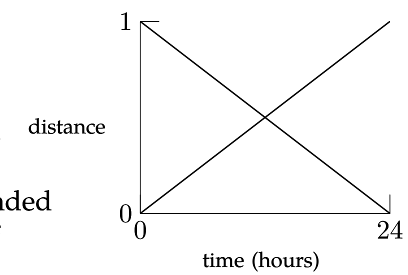

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

Maybe I'm missing something, but the graphs don't seem to correspond to the position curves described in the text. Not that it really matters, since the point is to show the intersection is hard to find without plotting them. Still, it's a bit of a detractor.

Well, the graphs are only showing the hour of the day, not the absolute time. It's really just time MOD 24 hours, which helps explain why they have to intersect.

This is not the best diagram for this problem, it is not intuitive from looking at the diagram and reading the problem how they are connected.

While I agree that it is certainly not immediately intuitive, all graphs require a moment's concentration to get your bearings. To decrease this time and considering the topic is on visual learning, a picture of a mountain on the y axis and a graphic of light and night on the x might help the reader jump to understanding an instant sooner. However, then this work would look too much like a pretty textbook, which detracts from its excellent rational and spartan grounding.

there aren't flat lines in the graphs...there should be if you're going to say there are...

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

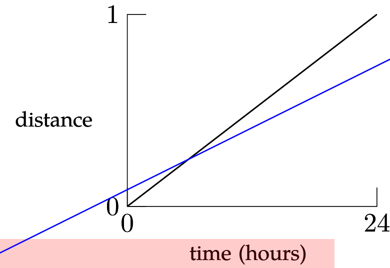

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

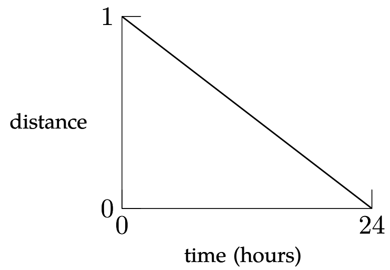

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

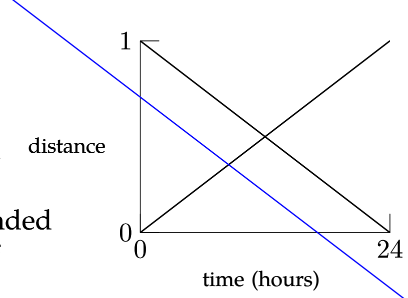

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

I'm not sure where the flat part of the diagram is? My guess is that you had a diagram with three parts: an upward part, a flat part, and a downward part, and then later decided to remove it and provide these diagrams instead without changing the text description.

I was also wondering about the "flat part of the diagram". However, it doesn't appear to be necessary to show it, since it is 24 hours and won't shift the schedule of d(t) relative to u(t). Thus, it's ok to just omit that part and overlap d(t) and u(t) exactly on top of each other since their time of day corresponds.

Yeah, this graph is sort of mod(24 hours), but without the flat part.

The 24-hour wait at the top seems like an oversimplification. Doesn't this still hold for any offset because time is mod(24) as someone pointed out? I pictured sliding the descending line left or right. I guess there's a trivial solution where you're only at the top and bottom at the same time of day

You could say that figuring out that the 24 hour long rest is arbitrary is part of separating the problem details from into the abstraction.

The flat part would be in the overall schedule, but in this overlay of graphs, we're only concerned with the time of day when you're at any given position, and the question says that you start going up and going down at the exact same time of day.

These diagrams don't really seem to correlate with what he is saying, since the slope appears to be constant. I understand that it doesn't matter, but it does make it confusing to say one thing and draw another. Either the text should be rewritten to say "constant speed" or the diagrams should be redrawn to match, including the level parts.

Now I remember what happened here. I wanted to make a diagram that corresponded to the text, i.e. with a flat part, but was in a rush to make a diagram, so I just put on in there that was "approximately right", with the intention of fixing it later. But I never did. Sorry about causing confusion with that.

but there's still no change in slope here, regardless of whether or not there's a flat portion

This makes sense, but the graph doesn't totally make that clear to me. Maybe there's just too many comment lines on the page that're distracting me.

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

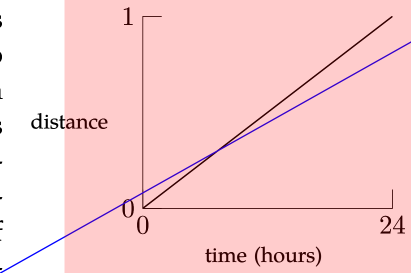

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

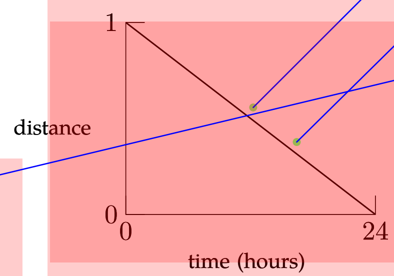

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

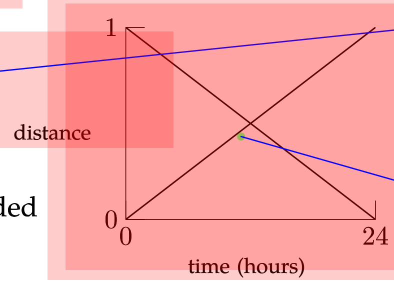

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

I don't see the flat line that represents resting...?

Me either I had the same question..

Same situation as with the first graph. Check the comments left on that one, it's sort of a typo.

This simplifies it nicely.

As the problem gets more complex, would it be possible to simplify it once to several diagrams, and then simplify those again to one or two diagrams?

If that's the case, then wouldn't it be easier to create a more simplified diagram?

The formatting of this page seems odd with random white space. It was rather confusing (particularly with the lack of proper diagrams a few paragraphs earlier). Would simply numbering the figures/diagrams, and referencing them in the text, make it easier to read?

i think latex takes care of all the formatting. i'm not sure how much control the author has over it.

It's probably because the diagram is connected to that paragraph, so although the words stop the paragraph continues until the end of the diagram.

I didn't think that you would be at the same point at the same time at first. And I wouldn't be able to express an argument any better than this last diagram.

I agree, the problem was very difficult to think about in my head but with these diagrams it is almost impossible to doubt the fact that you would be at the same place at the same time at some point.

While the graphs may not exactly match the description from the reading, the message trying to be sent seems really clear here.

Agreed, I really liked this example a lot.

Yea it's definitely very clear from this example how much easier it is to understand visuals quickly.

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

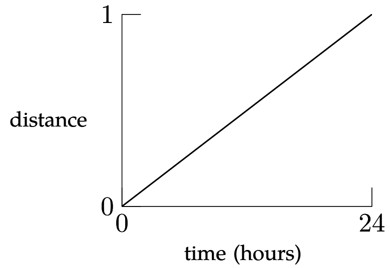

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

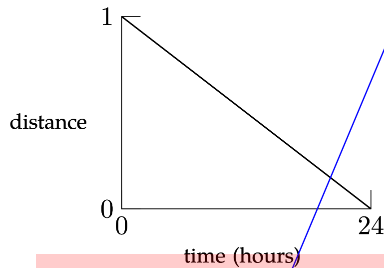

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

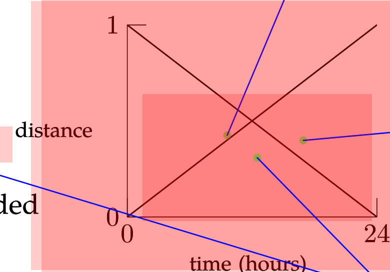

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

Usually I understand graphical arguments better than others, but for some reason these graphs don't convince me that there is always a point you reach at the same time going up and going down. Is there another proof of this?

Think of it as two people (this is how I originally solved it back in high school) one going up, one coming down. Will there be one point in time when they are on the same location on the mountain? Since they have to pass each other, the answer is yes. There is no difference here, except that you have to think of one person going both up and down.

Does the two people walking toward each other count as abstraction? Because in my head I'm picturing two things moving at each other.

I think there are other ways to describe this solution to this problem, but ultimately it reduces to this graph. The graph serves as the simplest form of the solution. I also believe that the curve does not necessarily have to be linear either on the way up or down, but either way the curves will intersect and you will be able to find a point at which you are the same distance at the same time of day.

The description of two people walking in opposite directions (given by the earlier commenter) really helped drive the point home, and ultimately made the graphs understandable for me. I think a description like this in the text might help solidify the abstraction presented by the graphs.

The main thing to notice is that the person has to walk down the mountain in the same time that he hiked up. That's what made the light bulb go off for me.

I think most people already solve problems visually like this in their head instead of the abstract way, but overall I think it is a great philosophy. A picture is always easier to understand than an abstract concept.

I'd say that when I think about a problem like this in my head I'm definitely using abstraction, even if i don't realize it.

Agreeing with previous readers, the diagrams should be changed, and just to reiterate your point (visually with diagrams!) you might consider including intersecting graphs with totally different schedules. Thus, one could look at three graphs and realize that yes, the paths will always intersect.

this quickly resolves any doubt or attempt to find a way to not be at the same place at the same time of day

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

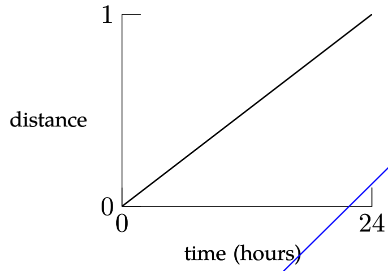

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

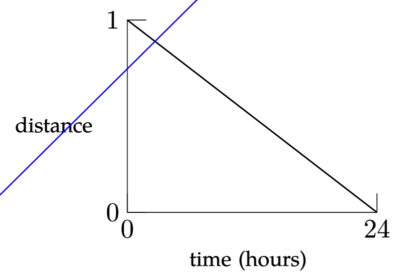

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

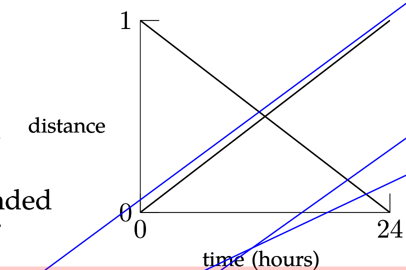

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

I still am not completely convinced. It seems to me that you should be able to walk in such a way that you won't be in the same spot at the same time of day. Like what if you walked fast on the way up and slow on the way down? Or maybe I am not fully understanding this diagram or the problem statement.

I had this same problem, but I think it's the specification that both up and down trips are over an entire 24-hour period, so even if you walk fast part of the time, the curves will still cross, even if it's just at some resting point.

If the diagrams don't help, another way to think about it is to imagine two people. One at the top one at the bottom, either walks at an arbitrary speed at one point they will pass each other provided they each reach their destination within 24 hours.

Thinking about the problem in terms of two people is really helpful..I think it would be worth including in the explanation of the diagram, otherwise it is somewhat difficult to understand why the solution is right.

I agree, that finished clearing it up for me.

Maybe include another possible set of diagrams (different slopes or flat parts in the middle of one or something) to really show the generality.

yeah I agree

I thought so

This is a good example to demonstrate the advantages of using methods that utilize visual processing. Could you give an example that is best solved with "perceptual processing"?

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

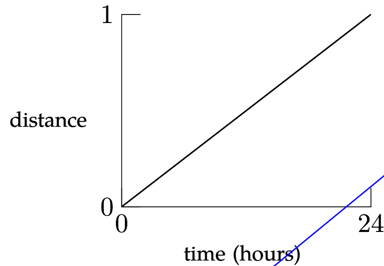

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

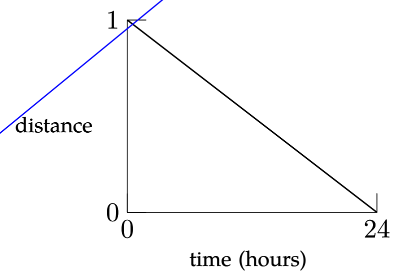

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

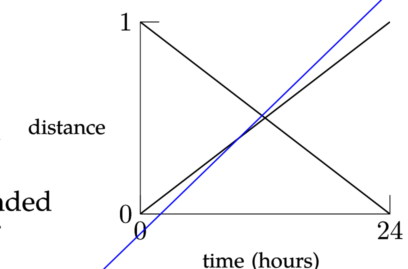

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

This is cool, my initial reaction (and as others have mentioned), is yes - my mind was thinking about the midway point, if you are hiking at the same speed, etc...you would be halfway at the same time. However, without the diagram and the abstraction, I'm not sure I would have realized quickly that it has to intersect once always (even if pace is different). I was thinking through other situations, and its definitely not easy without this method! This was a surprising for me and quite a useful method and example!!!

I agree, I still haven't been really sure of what exactly abstraction was and how it was useful, but this example sums it up really well. I wouldn't have been able to think through the problem as quickly without thinking about it this way.

I agree with this and think that once the graphs are updated it'll be even more definitive. This abstraction showed how useful it is to take the time to ignore the extraneous parts of a problem and that sometimes too much information can be detrimental

I think it would be interesting to illustrate a way of doing this problem without diagrams, just to compare the relative difficulties.

Yeah I agree, being able to see the advantage it brings would be pretty useful.

I don't know how you would... But maybe that's your point.

My initial reaction to reading the problem didn't think of this intersecting curves approach. Initially it seemed a little more thought intensive, but the explanation of the diagram helped simplify it greatly. The fact that this diagram approach seemed so simple seems like enough of a demonstration to show the benefits of using diagrams. Initial reactions from reading the text might be enough to serve as the "other" approach without diagrams.

I always thought graphing out a problem or making a picture added to the work. In this case, though, I may not have gotten the right answer independent of the diagrams.

this seems to be the only time you actually connect diagrams to abstraction

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

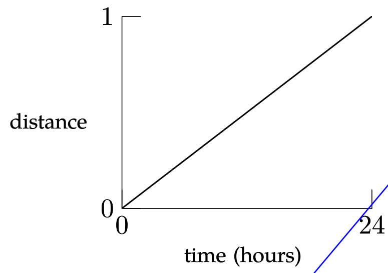

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

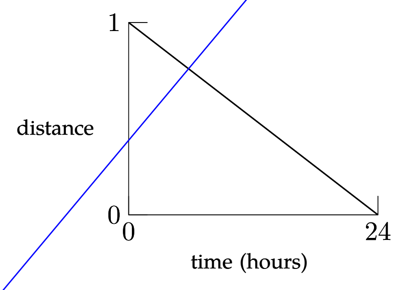

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

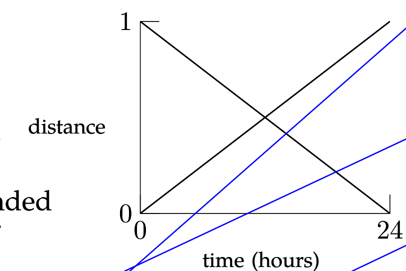

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

This problem is pretty intriguing to me (I've never seen it before). I've been trying to think of ways in which the paths wouldn't cross but I haven't been able to, it's pretty cool.

Yes, but I probably wouldn't have been able to come to that conclusion without the diagrams. Good example!

Without using this method of abstraction and drawing diagrams, this problem would have taken me much longer to do, and would have gotten a lot more complex

really? i thought the answer was apparent right away. although, i agree with others that another set of "drastic" plots would drive that point home.

I think having the diagram just makes the problem very simple to understand. Some may find this intuitive from the start, but the diagram essentially makes the problem easy to understand for people who didn't see it from the beginning.

We should have an example using Venn Diagrams as well

I really liked this explanation because without it I would have come up with myriad hypothetical alternatives where it didn't hold true and maintained that I was right.

good summary

I thought the answer 'Yes' was intuitive. After I read the problem, it just seemed like 'This has to be yes, right'. I don't need the need for this example, but I do agree with an above comment, I'd like to see the diagram.

I really liked this example- I actually understand what you are talking about now

straightforward and easy to understand reading.

This final point would be more clear if you used graphs that were nonlinear as well (start out fast, slow down, etc)

This reading was very easy to follow.

Except for the mismatch between the diagrams and the text, I think that this section is overall very clear and well written. I honestly have no comments (that would not almost be verbatim to what was already said) except the one I made above.

the question is: Must $u(t) = d(t)$ for some value of t? Or can you choose $u(t)$ and $d(t)$ to avoid the equality for all values of t?

With these abstractions, the question is cleaner, but it is not yet easy enough to answer. A diagrammatic representation makes the answer more obvious. Here is a diagram illustrating an upward schedule. Distance is measured from the bottom of the mountain (0) to the top of the mountain (1). According to the indicated schedule, you walked fast (the initial slope), rested (the flat part), and then walked to the top.

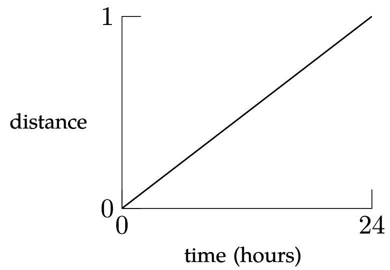

Here is a diagram illustrating a downward schedule. On this schedule, you rested (initial flat line), walked fast, then walked slowly to the bottom.

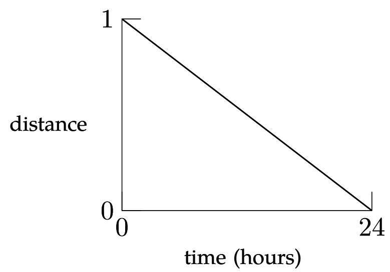

And this diagram shows the upward and downward schedules on the same diagram. Something interesting happens: The curves intersect! The intersection point gives the time and location where the upward and downward schedules landed on the same point at the same time of day (but on separate days).

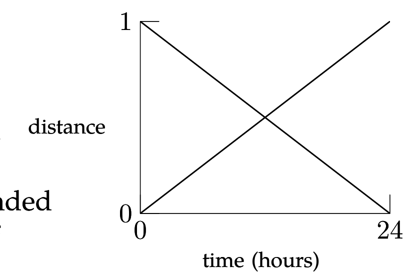

Furthermore, the diagram shows that this pattern is general. No matter what schedules you choose, the upward and downward paths must cross. So the answer to the question is 'Yes, there is always a point that you reached the same time on the upward and downward journeys'– an answer hard to reach without abstracting away all the unessential details to make a diagram.

looking back, I think i might have been doing a lot of proof reading around the time of this (& the previous) reading

2.2 Recursion

Abstraction involves making reusable modules, ones that can be used for solving other problems. The special case of abstraction where the other problem is a version of the original problem is known as recursion. The term is most common in computing, but recursion is broader than just a computational method – as our first example illustrates.

2.2.1 Coin-flip game

The first example is the following game.

> Two people take turns flipping a (fair) coin. Whoever first turns over heads wins. What is the probability that the first player wins?

As a first approach to finding the probability, get a feel for the game by playing it. Here is one iteration of the game resulting from using a real coin:

```
TH
```

The first player missed the chance to win by tossing tails; but the second player won by tossing heads. Playing many such games may suggest a pattern to what happens or suggest how to compute the probability.

However, playing many games by flipping a real coin becomes tedious. A computer can instead simulate the games using pseudorandom numbers as a substitute for a real coin. Here are several runs produced by a computer program – namely, by a Python script `coin-game.py`. Each line begins with 1 or 2 to indicate which player won the game; then the line gives the coin tosses that resulted in the win.

```
2 TH
2 TH
1 H
2 TH
1 TTH
2 TTTH
2 TH
1 H
1 H
1 H
```


From the game's description it doesn't seem like half the time is a plausible answer. Player 1 definitely has a higer chance of winning...

where did this 1.58 come from?

I was a little confused at this being the school method. Took me a while to realize I do actually calculation complex multiplication like this but vertically. Maybe it would be easier to show it the vertical way. Also every multiplication we do as humans is the recursion or multiplication since build on what we know via our memorization of basic multiplication (ie. 3*5=15 so 50*3=150).

I really don't understand how this was derived. Before I checked it, it seemed like you just split up the numbers and adding/multiplied them around several times. Does this work for any complex multiplication or is the a special case? Please review in class.

I guess this is similar to when you multiple vertically because we do cross multiply digits and them sum the results up. It may just seem cooler when its explicitly written out and shown horizontally.

I don't understand the subtraction portion.

I think the examples in the chapter did a good job in explaining what recursion is because initially the definition just sounded like abstraction and divide and conquer.

2.2 Recursion

Abstraction involves making reusable modules, ones that can be used for solving other problems. The special case of abstraction where the other problem is a version of the original problem is known as recursion. The term is most common in computing, but recursion is broader than just a computational method – as our first example illustrates.

2.2.1 Coin-flip game

The first example is the following game.

> Two people take turns flipping a (fair) coin. Whoever first turns over heads wins. What is the probability that the first player wins?

As a first approach to finding the probability, get a feel for the game by playing it. Here is one iteration of the game resulting from using a real coin:

```
TH
```

The first player missed the chance to win by tossing tails; but the second player won by tossing heads. Playing many such games may suggest a pattern to what happens or suggest how to compute the probability.

However, playing many games by flipping a real coin becomes tedious. A computer can instead simulate the games using pseudorandom numbers as a substitute for a real coin. Here are several runs produced by a computer program – namely, by a Python script `coin-game.py`. Each line begins with 1 or 2 to indicate which player won the game; then the line gives the coin tosses that resulted in the win.

```
2 TH
2 TH
1 H
2 TH
1 TTH
2 TTTH
2 TH
1 H
1 H
1 H
```

Read Section 2.2 (memo due Sunday at 10pm).

This is kind of awkward wording. I had to read the sentence a couple of times to figure out what you mean.

yea, when I think about it I usually think of it as a problem inside a problem, like those wooden dolls that open up and another little replica is inside and it keeps going until you get a tiny doll. That example works for me.

It might be useful to include one or two sentences in the beginning explaining how the coin-flip game applies to recursion. While it is reasonably clear by the end of the section, a few sentences might do a lot to orient the reader.

I disagree, I think the way this section is structured is perfect and does a good job of letting the reader discover the idea of recursion on their own without holding their hand and telling them exactly how the example relates in the beginning.

I disagree and believe it is better this way. It makes the reader think a bit more throughout the section and is better for learning.

I can appreciate both sides of this discussion. There's always a tension between presenting results as in a handbook (i.e. for those who already know a result, and discussing results so as to help readers best learn them. Mostly I've chosen the learning side, but I sometimes wish there could be a two-layered book: Once you read it and learn everything in it, it turns into a handbook!

2.2 Recursion

Abstraction involves making reusable modules, ones that can be used for solving other problems. The special case of abstraction where the other problem is a version of the original problem is known as recursion. The term is most common in computing, but recursion is broader than just a computational method – as our first example illustrates.

2.2.1 Coin-flip game

The first example is the following game.

> Two people take turns flipping a (fair) coin. Whoever first turns over heads wins. What is the probability that the first player wins?

As a first approach to finding the probability, get a feel for the game by playing it. Here is one iteration of the game resulting from using a real coin:

```
TH
```

The first player missed the chance to win by tossing tails; but the second player won by tossing heads. Playing many such games may suggest a pattern to what happens or suggest how to compute the probability.

However, playing many games by flipping a real coin becomes tedious. A computer can instead simulate the games using pseudorandom numbers as a substitute for a real coin. Here are several runs produced by a computer program – namely, by a Python script `coin-game.py`. Each line begins with 1 or 2 to indicate which player won the game; then the line gives the coin tosses that resulted in the win.

```
2 TH
2 TH
1 H
2 TH
1 TTH
2 TTTH
2 TH
1 H
1 H
1 H
```

It is worthwhile to note that the condition is based on whether the first player wins. Something I took to mean the first "toss", but it actually refers to the *player* who tosses first. So the condition will still be true if the player wins on the third toss (or fifth, etc.).

I don't really understand what you found confusing in this little paragraph. Is your confusion with "first" meaning toss or player actually farther down?

On first read I too thought it meant "what is the liklihood that the first player wins on the first toss (1/2)" rather than the correct reading of "what is the liklihood that the player who goes first wins at all?" I think its just a priming thing.

I agree - I got a little confused about it the first time, interpreting it as "first player winning the first toss" but after reading it again it's fine.

Yeah I think what this means here is that coins are flipped until someone gets heads. That player is the winner. Then what is the probability the first player is that winner?

just as a counter-perspective, i really didn't have any problems understanding this game set up and all and thought it was very clear.

just as a counter-perspective, i really didn't have any problems understanding this game set up and all and thought it was very clear.

just as a counter-perspective, i really didn't have any problems understanding this game set up and all and thought it was very clear.

just as a counter-perspective, i really didn't have any problems understanding this game set up and all and thought it was very clear.

I also thought it was clearly stated.

this was a little confusing to me. then i realized what it represented.

Agreed. Although I think we've all seen this notation before, if you wanted this book to be accessible by non-math oriented people, then you should probably introduce it as Tails, Heads before using TH.

Or, it could be stated that T means tails and H means heads, so that it is stated for the rest of the document.

2.2 Recursion

Abstraction involves making reusable modules, ones that can be used for solving other problems. The special case of abstraction where the other problem is a version of the original problem is known as recursion. The term is most common in computing, but recursion is broader than just a computational method – as our first example illustrates.

2.2.1 Coin-flip game

The first example is the following game.

> Two people take turns flipping a (fair) coin. Whoever first turns over heads wins. What is the probability that the first player wins?

As a first approach to finding the probability, get a feel for the game by playing it. Here is one iteration of the game resulting from using a real coin:

```
TH
```

The first player missed the chance to win by tossing tails; but the second player won by tossing heads. Playing many such games may suggest a pattern to what happens or suggest how to compute the probability.

However, playing many games by flipping a real coin becomes tedious. A computer can instead simulate the games using pseudorandom numbers as a substitute for a real coin. Here are several runs produced by a computer program – namely, by a Python script `coin-game.py`. Each line begins with 1 or 2 to indicate which player won the game; then the line gives the coin tosses that resulted in the win.

```
2 TH
2 TH
1 H
2 TH
1 TTH
2 TTTH
2 TH
1 H
1 H
1 H
```

Does the first player always have the first flip?

I think so in the case of this example.

It seems that the first player is *defined* by the first toss.

I agree it would be a different probability if he didn't.

Exactly – it doesn't matter whether Alice or Bob wins, just whether the first player (ie the one who tosses first) wins.

this would be the perfect place to have an aside as to what a pseudo-random number is...as a side note, I really like that you actually [and correctly] call it a psuedo-random number [as opposed to a random number]

how is pseudorandom different from random?

pseudorandom means it approximates the properties of random numbers

Taken from wikipedia: "A pseudorandom number generator (PRNG) is an algorithm for generating a sequence of numbers that approximates the properties of random numbers. The sequence is not truly random in that it is completely determined by a relatively small set of initial values, called the PRNG's state."

Basically, the PRNG is given an initial number (or set of numbers), and by performing certain operations on that number it arrives at a sequence of "random numbers." The pseudo part comes because knowledge of the initial numbers determines the rest.

What about just using math to calculate the probability? Computer simulations provide an estimate of the expected probability, but doesn't really add to the logic behind the theory, right?

Yes, especially since the simulation is technically not entirely random. How bout the course 2 students make a coin flipping robot?

There are also many calculator games which do this. I realize its unrelated, but in high school they made us put them on our calculators to do probability simulations

I think the calculator games are also "pseudorandom number generators". The idea is that any computational device, like a computer or calculator, can't truly pick a random number, and can only simulate a random choice. So I think the calculator games you mentioned are the same as the computer program described here.

2.2 Recursion

Abstraction involves making reusable modules, ones that can be used for solving other problems. The special case of abstraction where the other problem is a version of the original problem is known as recursion. The term is most common in computing, but recursion is broader than just a computational method – as our first example illustrates.

2.2.1 Coin-flip game

The first example is the following game.

> Two people take turns flipping a (fair) coin. Whoever first turns over heads wins. What is the probability that the first player wins?

As a first approach to finding the probability, get a feel for the game by playing it. Here is one iteration of the game resulting from using a real coin:

```
TH
```

The first player missed the chance to win by tossing tails; but the second player won by tossing heads. Playing many such games may suggest a pattern to what happens or suggest how to compute the probability.

However, playing many games by flipping a real coin becomes tedious. A computer can instead simulate the games using pseudorandom numbers as a substitute for a real coin. Here are several runs produced by a computer program – namely, by a Python script `coin-game.py`. Each line begins with 1 or 2 to indicate which player won the game; then the line gives the coin tosses that resulted in the win.

```
2 TH
2 TH
1 H
2 TH
1 TTH
2 TTTH
2 TH
1 H
1 H
1 H
```

Is there any way you could include the script in an appendix or something? I would be interested in seeing it.

I would also be interested in seeing the script. Also, unless the script is included, I'm not sure why it is relevant to include the name of the script.

Sure. I've just posted it on the course website (in the "data/scripts" section).

Awesome, which version of Python is this written in? I believe we used 2.4 for 6.00.

is it really necessary to have this detail in the text then?

I think it would be useful to have an appendix devoted to all the code in the book and maybe a brief description, or if there's code that explains a concept well that does not fit in the flow of the book proper. It would help with understanding for the code-inclined.

I agree – a code appendix would be very nice, and could easily be ignored by someone who wasn't interested in it.

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

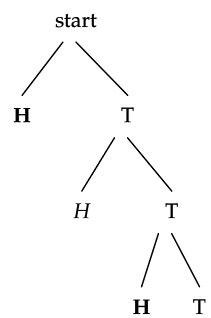

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \qquad (2.1)$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, p = 2/3, as conjectured from the simulation.

shouldn't player 1 have won 50% of the time on the first try? maybe that'll hold true when the simulation is run more times.

what we're calculating isn't who wins the first time, but who ultimately wins. so if the first player didn't win the first time, we go to the second player, and so on and so forth, until one of the two players gets a head and wins the game.

I think you may have been caught in the confusion of the sentence describing the win condition. As the comment before this says, we're calculating how often player 1 is the first to win, not when he wins on the first try.

"Reasonable" conclusion? I don't know that many people would make this assumption.

You might if you knew nothing about the game. If you were only told that there was a game that Player 1 won 5 out of 10 times (without being told how the game worked), you would guess it was fair, or at least close to fair.

Even if I didn't know anything, I'd still be suspicious. 10 throws is certainly better than just 2, but not enough to make any claims of fairness.

Yes, but if someone asked you "which player would you like to be" (again, without explaining the game), you wouldn't have any reason to prefer one or the other, and moreover you wouldn't have any reason to suppose that you /should/ have a preference.

i just don't think it's a big deal. the phrase is fine.

yeah i don't know what you guys are arguing about. the point he's trying to make is obvious.

Hearing that it is fair and looking at the example and who won are totally different. It is obvious player one has an advantage because they have to lose before 2 can win. So how could it be fair?

I think the point of this was just that if you didn't understand the problem the observed data would show you that it was fair. Not that the probability shows that the game is fair.

I agree. If you just think of it as a black box, the output makes it look like a fair game from the outside.

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

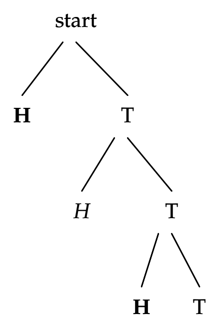

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \quad (2.1)$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

I think that a side-bar would be useful here...explain what, exactly, you mean by fair. I mean, I get it, but i'm not sure it's something that would be intuitive to [most] everyone that might read the book

I would think that this might give a slight advantage to the person who goes first

Minor edit, but this line seems a little too wordy on first read through, enough to confuse me (I thought the placement of 'based' was a typo)

A comma after strongly or striking "as it is" might make it clearer.

You are right. Leaving out helpful or even all commas is a confusing habit that I picked up in England (I lived there 30% of my life). [Whoops, I did it again in the preceding sentence.] The habit returns like a retrovirus when I least expect it.

I think the sentence would be clearer if it were slightly reordered to read "the conclusion cannot be believed too strongly, as it is based on only 10 games" (comma debatable).

I would consider jsut moving the word 'based' after 'as it is'. That seems more grammatical in my opinion.

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

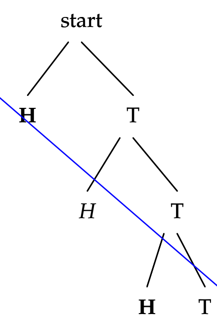

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \tag{2.1}$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

Is there a way of determining exactly how many games would be necessary? A number that would lead to a 'confidant' result?

I have often wondered about what makes a probability legit. For example, in a lot of papers I read, simulations or experiments will show that "98.9%" of the time something they want happens. However, I was talking to someone who does a lot more probability than I do (in biology), and he said he would only trust something that is 99.9999% or along those lines. Does it just depend on the thing that you are looking at? Such that if you are dealing with living things, that you must have a larger sample to account for variability?

I think this is always a hard question to answer. From what I know about probability and statistics (which isn't much), there are many ways you can go about doing this. t-tests, chi square tests, etc. to guage for how accurate certain values are; and you could always look at standard deviation and z scores as a measure of how far off measurements are from each other–larger sample leads to smaller standard deviation

Well I'm not entirely sure, but I believe you could use the law of large numbers to solve this. The LLN states the lim(n->inf)Prob(|Xn-u|<e)=1 for Xn average of the trials, u expected value of the outcome, and e error. By choosing a small e we can produce a "confident" answer and figure out the number of trials n.

That's a very interesting question. I've been looking for examples for the (only half-written) chapter on "Probabilistic reasoning". After reading your question, I think I'll use this coin-game simulation as one example. We'll figure out how confident one can be about p as a function of the number of games in the simulation.

(For those who already have familiarity with statistical inference: The analysis will be Bayesian.)

central limit theorem says we must approach the true result with more number of trials

I think the theorem says it works fine above n = 30

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

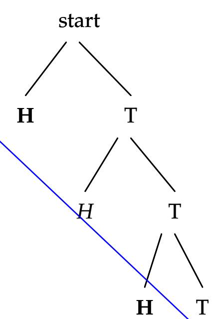

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \tag{2.1}$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

Why don't you just build it into the coin.game code?

Abstraction!

Haha, exactly! That's sort of the same idea behind creating functions in programs. When you make a function, you define it and allow it to be called later, which not only saves you time (and redudant typing) but also gives you the ability to access and combine these 'modules'.

It would be significantly faster though...

In case anyone else has to look this up... wc is the command for word count

i don't bother to look them up... i just skim over all the programming and go to the sections where he explains the actual estimation technique.

That's probably where the dividing line is between the course 6 and course 2 majors...

I hope the line gets blurred by the end of the course!

Not extremely important, but "most recent" is somewhat ambiguous here. Should I take it to mean the most recent in terms of when this book was written? Perhaps use "one" instead?

My guess would be that "most recent" refers to one, non-specific pipeline that was run at the time this section was written.

This makes sense as the idea of the inclusion of pseudorandom would cause different experimental results, but the wording is rather awkward. The important part is the 68 to 32 split between the two players and I think this should be the focus.

I think that the wording "most recent" is a little off. I understand that it is meant to place the time of the results, but I'm not sure that's necessary in a book. Maybe list a set of three (or so) results?

Perhaps you could also discuss how large n should be for such results to be significant.

the first player has an inherent advantage, but this doesn't mean it's not fair, right?

"fair" would imply that each player has an equal (1/2) probability of winning each game. If the game is set up such that one player is more likely to win, it doesn't meet the definition of fair.

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

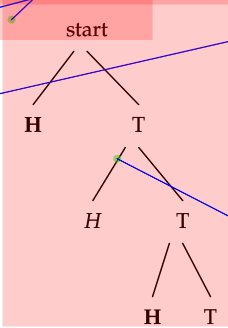

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \tag{2.1}$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, p = 2/3, as conjectured from the simulation.

Doesn't mean much- In the 10 iterations, it is still more probable it had a ratio closer to 2/3

This kinda makes sense because the first player has to lose for the second player to even have a chance to win. A one time coin toss would be fair but this is a little different.

This tree made it much easier of understanding the premise for the game

Maybe I'm confused now – I thought the premise was you and another person alternate flipping coins until someone gets heads, right?

yes, that is the premise. I think that the tree does nothing for better explaining the game, but it does improve the understanding of the probability calculations

I don't quite understand how this tree is a representation of the game- what happens if the coin has to be flipped three time (tails, tails, head) where does that come in?

So if the boldface represents the first player, I assume the italics represent the second player?

I think that the only italic is the second H down the tree...and i'm pretty sure that's a typo...i do know it should be bold

No, I don't think it's a typo. It should indeed represent the second player winning because the pattern that leads up to it is TH, which means that the first player threw tails, and then the second player threw heads and won.

We used this method a lot in high school statistics, but they never introduced the concept of recursion with it.

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

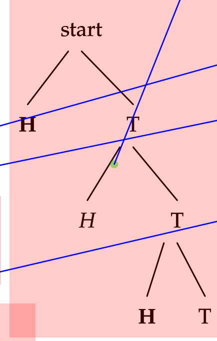

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \qquad (2.1)$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

The tree makes it so easy to understand now

I feel this would be the best way to find the anser to the problem

It might be helpful to add a line next to each toss showing which player is tossing. Also, it's a minor point, but the tree doesn't show the TTTTH scenario mentioned in the text.

I agree. The tree makes the answer seem obvious,

i know that this will be better explained later on in this section, but when i first saw this tree, my first reaction was "isn't that from divide and conquer? aren't we supposed to be on abstraction now?" did anyone else do this?

I think maybe the point of using the tree is to show that trees have many useful applications in different areas and in different problem solving methods.... divide and conquer, abstraction, etc.

yeah so if we had this, why did we run a program in the first place?

How did you get these probabilities?

I've never seen this result before and found it pretty interesting and surprising.

I think that this is completely clear, but for someone who doesn't have any probability (given that probability isn't a pre-req for this course), you may want to explain how that sum comes around.

I agree. This conclusion definitely could be explained in greater detail. I'm confused on why the probability doesn't remain 1/2 the whole time or why it doesn't go from 1/2 to 1/4

because the second one P(TTH) = 1/2 * 1/2*1/2 = 1/8, P(TTTTH) = 1/2 ^ 5 = 1/32. 1/4 (TH) is not what we want, because (TH) means the second person gets a head, but we want to calculate the probability of the 1st person gets the head

I agree... clearly explaining that you are calculating the probability that only the 1st, 3rd, 5th etc (only taking odd terms) would be helpful

or maybe just include exponents. p = 1/2 + 1/2^3 + 1/2^5 ...

is this the sum of 1/2*4^n?

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

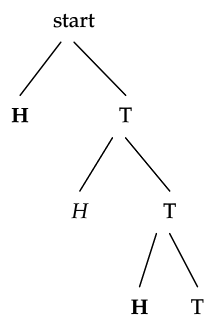

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \tag{2.1}$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and $p/2$ add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

I always forget in probability (I never ever have to think about it!) when to add, or multiply, or use exponents. Does anyone have a good way of remembering, so that when I do need to use it once every year or so, I can do it right?

I think if there are n ways to reach the same result, in this case, n ways to win, then you add up the probabilities. however, if a series of things need to happen in order to achieve a certain result, then you multiply, this is just one way to understand it.

In probability, "and" corresponds to multiplication and "or" corresponds to addition of probabilities. in this case, the probability of the first player winning is if the first coin is heads OR if the third coin is heads OR if the fifth coin is heads; thus, the probabilities are added

Thank you!

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

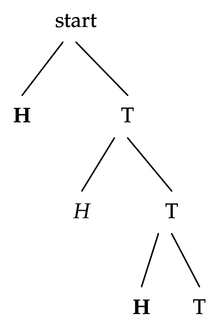

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \quad (2.1)$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

What formula is that? I forget...

I agree, the formula should be included.

sum(a*r^n) = a/(1-r) over n=0 to +inf, for |r|<1. Here, r=1/4, and a=1/2. (1/2)/(1-1/4)=2/3

Thank you for the equation! I agree it should probably be included in the reading.

I don't think it has to be. I consider a good textbook one that gives you the basic ideas, and then encourages you to go out and explore things for yourself. Besides, it's pretty fun to figure out the formula yourself if you have forgotten.

oh yeah...fun...if you're a masochist.

Fun exploration is trying out the new knowledge you gained, and playing around with it. Tedious, unnecessary 'exploration' is google searching "geometric series formula" to get an equation that would have taken up less space than the phrase "a familiar formula".

Even if you wanted figure out the sum for yourself, it'd be nice to know if you got it right.

Although I see how it might be nice to see the formula it is somewhat besides the point of this chapter. We are using abstraction in a lot of these examples to prove a point, and going into too much detail in every example just detracts from the actual point.

In the next paragraph, I explain an alternative method that you can use instead. Then you can use the alternative method to derive the geometric-series formula.

when can you use this formula- not for all series? are there a few formulas we should memorize for these types of series?

Equation or problem 2.1? This seems to refer to the equation above, not the problem at the end of the section.

It is indeed Problem 2.1, whose solution gives the familiar formula needed for Equation 2.1 (though until one solves the problem, it's not so evident).

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

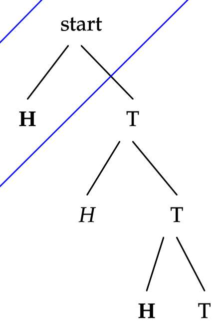

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \quad (2.1)$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

so in this instance, the tree method and recursion method are used interchangeably on this problem. i am a more visual person and find the tree more intuitive than the written unix code. is there a rule of thumb for when one is more appropriate than another?

If you consider who wins as part of the tree structure, then it's really sort of 'reversed with respect to player number' one level down... and hence the next result. If it were truly actually repeated one level down, then the game would just be Player 1 flipping a coin until he won, right?

This is clever. I would have easily realized that p1+p2=1 but i would not have so easily realized that p2=1/2*p1. (Where p1=probability of 1 winning and p2=probability of 2 winning)

I agree- this is really awesome. I would never have come up with this.

I agree–this is really clever. The tree diagram helps a lot with visualizing the recursion–every time there's a tails, the same process happens again, and here, for the second player to win, the first player first has to get tails, then the second player has to get heads–1/2 the chance of the first player winning since theres that extra condition of the first player getting a tails

The probability a player wins a toss starts at 1/2. However the second person only gets to go if the first person losses (half the time). Only two people can win so p+p/2 = 1. and now we know p, which is percentage for first player.

This is pretty cool. I wouldn't have guessed this either at first, but after looking at the tree, it makes sense.

I didn't quite get it at first, but the class comments really cleared it up!

Yeah, this is amazingly simple and I didn't get it at first either. Pretty cool.

I am still kinda not completely convinced about all this. It still seems kinda intuitive to me that if this is a recursive example then the probability should be the same every time and thus should add to equal 1. The class comments do help and if I am correct then the reason it is not 1 is because player 1 gets the first flip? I don't really see how that would affect the probability though. (shouldn't it not matter who goes first?)

Would be even more clear if you added these probabilities to the tree and show how it iterates

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

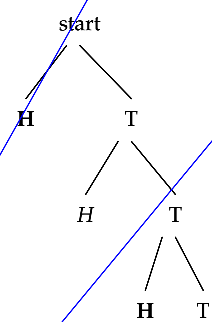

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \quad (2.1)$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and p/2 add to unity. Therefore, p = 2/3, as conjectured from the simulation.

This analysis is much easier for me to follow.

I agree, I like this analysis, although it seems a little less intuitive, which makes it a nice addition in the text.

Agreed, this seems very simple and it's very interesting that it really is this intuitive.

I think the effect is even greater given the contrast with the more complicated math of the previous section

I have to agree that it was much easier to follow this when compared to the previous section. It's clear and simple, without much outside knowledge needed.

I also think so, I actually started thinking about it this way first.

Where p is the probability that the first player loses the first round? Because if that's the case how do you get p = 2/3...

I'm not sure I follow what's happening in this paragraph.

Prob that P1 wins = p, Prob that P2 wins = 1/2*p. They add to unity, so p + 1/2*p = 1. This reduces to 3/2*p=1, or p=2/3

i don't understand either. all variables should be really clear, and it's easier to understand math when it's not inside of a paragraph

This seems pretty obvious to me. If the first player always starts, he's going to win more. Am I not getting something?

There is a disconnect for me on this step. I get it from a theoretical standpoint, however, it still seems that the probability should be 50/50.

p = 2/3 is really counter-intuitive. One first pass, I thought it would be 1/2 by symmetry, but I guess since player 1 flips first, it breaks the symmetry argument. I like the p + p/2 = 1 argument a lot though.

I'm still misunderstanding this part.

It might be helpful to just write out explicitly the two equations you use here to solve, as a summary:

p1 + p2 = 1 p2 = (1/2)*p1

In these ten iterations, each player won five times. A reasonable conclusion, is that the game is fair: Each player has an equal chance to win. However, the conclusion cannot be believed too strongly based as it is on only 10 games.

Let's try 100 games. With only 10 games, one can quickly count the number of wins by each player by scanning the line beginnings. But rather than repeating the process with 100 lines, here is a UNIX pipeline to do the work:

```
coin-game.py | head -100 | grep 1 | wc -l
```

Each run of this pipeline, because `coin-game` uses different pseudorandom numbers each time, produces a different total. The most recent invocation produced 68: In other words, player 1 won 68 times and player 2 won 32 times. The probability of player 1's winning now seems closer to 2/3 than to 1/2.

To find the exact value, first diagram the game as a tree. Each horizontal layer contains H and T, and represents one flip. The game ends at the leaves, when one player has tossed heads. The boldface H's show the leaves where the first player wins, e.g. H, TTH, or TTTTH. The probabilities of each winning way are, respectively, 1/2, 1/8, and 1/32. The infinite sum of these probabilities is the probability p of the first player winning:

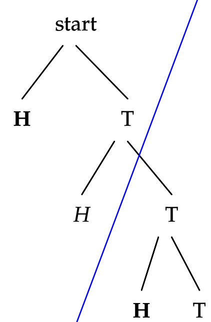

$$p = \frac{1}{2} + \frac{1}{8} + \frac{1}{32} + \cdots. \tag{2.1}$$

This series can be summed using a familiar formula.

However, a more enjoyable analysis – which can explain the formula (Problem 2.1) – comes from noticing the presence of recursion: The tree repeats its structure one level down. That is, if the first player tosses tails, which happens with probability 1/2, then the second player starts the game as if he or she were the first player. Therefore, the second player wins the game with probability 1/2 times p (the factor of 1/2 is from the probability that the first player tosses tails). Because one of the two players must win, the two winning probabilities p and $p/2$ add to unity. Therefore, $p = 2/3$, as conjectured from the simulation.

I really like this example. The idea of recursion is a very powerful way to analyze the game's probability. However I think the analysis could be improved by expanding the last paragraph, and possibly including a few more drawings.

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \quad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \quad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \quad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \quad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \quad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \;\times\; 2|7 = 3 \times 2 \;\Big|\; 3 \times 7 + 5 \times 2 \;\Big|\; 5 \times 7. \quad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

This seems a bit out of place...and I'm not sure I know enough to do this!

How does this relate to the examples above? And I think you need to add the condition that abs(r) < 1 for this question to have any kind of meaning.

On the previous page (line 2.1) we compute the probability of the first player winning by adding his chance of winning throughout the inifinite tree. However, there are an infinite number of terms to add, but since the terms are powers of 1/4, (p = (1/2)(1+1/4+1/4^2+...), the formula simply helps you get the answer to the sum of the infinite terms.

Yeah I agree, it should be noted that the condition of abs(r) < 1.

Let Z = 1 + r + r^2 + r^3 + ...

Z = 1 + r(1 + r + r^2 + ...) Z = 1 + r(Z) Z - rZ = 1 Z(1-r) = 1 Z = 1 / (1-r)

Wow, that was helpful! And the simplification from step 1 to 2 is a great example of abstraction. Maybe this could be included as an example before we are prompted with a problem to solve on our own?

You do know enough, but it requires thinking about it in an interesting way (and it's a problem on HW 2).

Problem 2.1 Summing a series using abstraction

Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots . \quad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \quad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \quad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \quad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \quad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \ \times \ 2|7 = 3 \times 2 \ \big| \ 3 \times 7 + 5 \times 2 \ \big| \ 5 \times 7. \quad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

wouldn't this just be infinite if |r|>=1?

yeah.. it's conventional to state a |r|<1 bound, although the answer could include cases, |r|<1 and |r|>=1.

Maybe it should be stated in the problem prompt. Although I hadn't actually considered the above case, I think it would be nice to have the specifcs.i

Maybe it should be stated in the problem prompt. Although I hadn't actually considered the above case, I think it would be nice to have the specifcs.

I think the purpose of this question is not to be mathematically rigorous here, but rather to find interesting, pictorial or abstractional ways to find the sum of this series

but the abstraction falls apart for |r|>=1, and bounds are usually VERY important. If you perform the abstraction, get an equation, and forget to determine and include the bounds, then can run into big trouble. For r=2, in this example, your equation would yield Sum = -1. In this case a sense check can quickly show that something is wrong, but not every situation is so clear.

On the other hand, there are whole areas of mathematics and physics where one just goes ahead and sums the series, worrying about the rigor later (if ever). Quantum electrodynamics is a good example.

Given the abundance of rigor in most education, it's worthwhile to suspend, for a while, the quest for it.

I agree about focusing on other ways to solve it as opposed to mathematical rigor, however, since recursion was just introduced with only one example, i feel that this problem is a little advance for people to begin drawing on recursion principles. At first look I didn't even really understand how recursion applied until i saw it explained in the comments.

I disagree – I think that not stating that sort of condition in the prompt is appropriate in this circumstance. It's important for us to be able to distinguish between the cases in which the series might converge or not, and a real application might not give us all the restrictions we would like.

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \quad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \quad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \quad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \quad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \quad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \times 2|7 = 3 \times 2 \;\Big|\; 3 \times 7 + 5 \times 2 \;\Big|\; 5 \times 7. \quad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

How is this useful for us if we're learning to approximate in order to not use a machine?

Think of the goal here not only as learning the approximate (I admit, the course title is misleading), but as learning to use the reasoning tools, like abstraction, across many, many fields. So, here you will see how abstraction leads to recursion which leads to a clever multiplication algorithm.

i'm not sure i understand how this is abstraction though

It's abstraction in that it takes a higher level routine and calls itself to accomplish a smaller task.

What is the standard school method?

But we are humans. So how does this method help us humans approximate? Wouldn't we just use the 1 few or 10 method and do it in a few seconds instead?

The one/few/10 method is good for approximations. I think this is for when we need a more exact answer.

Or for when you want to teach a computer how to do large calculations. (I really need to change the title of the course so that I can indicate that the course is not just approximation...)

Just curious- Pi is an irrational number...so what are computers calculating mathematically when they calculate pi to billions of digits?

"Although practically a physicist needs only 39 digits of Pi to make a circle the size of the observable universe accurate to one atom of hydrogen, the number itself as a mathematical curiosity has created many challenges in different fields."

http://en.wikipedia.org/wiki/Pi#Computation_in_the_computer_age

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \tag{2.2}$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \tag{2.3}$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \tag{2.4}$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \tag{2.5}$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \tag{2.6}$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \;\times\; 2|7 = 3 \times 2 \;\big|\; 3 \times 7 + 5 \times 2 \;\big|\; 5 \times 7. \tag{2.7}$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

I have no idea what this is. Maybe replace it with a more common knowledge example?

I agree that anyone outside of course 6 wouldn't have any idea what public key cryptography is. But for anyone still confused, wikipedia has a nice explanation: http://en.wikipedia.org/wik key_cryptography

I think public-key cryptography is one of the most important concepts that has emerged in recent decades. Even in its simplest form, it should be taught in any class (not just in course 6) for the simplicity and logic of the idea.

I think this is a good example in an electronic format where you can google something you don't know. Maybe it's not a common example, but he's trying to accommodate multiple audiences as can be seen by the title of the class.

Is this sort of like divide and conquer, but then apply the same operation to each smaller part(recursive step)?

I don't understand this way of computing, it seems somewhat difficult and extensive. Also, how is this useful for us as approximaters?

I've never heard of this

Agreed. Is it just how people are taught to multiply by hand?

I too have never been taught the "school method" shown here, but if the example is only meant to illustrate how multiplication can be tedious without abstraction then it still seems to serve its purpose even if you haven't seen that particular method.

I don't think I was taught how to multiply with this. It was more of a "this is what these numbers can breakdown into."

I'm also confused what is meant by the "school method". The only method I was ever taught in school was where you lined the numbers up vertically and multiplied each digit through, etc.

I think the idea, although it is not exactly like how we solved problems in school, is that we are isolating the various parts of the problem. Aside from the tricks of grouping these numbers, we are still multiplying the same way as always.

Is the school method something you are making up right now, or is it an establishes technique? I have never herad of it before

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \qquad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \qquad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \qquad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \qquad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \qquad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \;\times\; 2|7 = 3 \times 2 \;\Big|\; 3 \times 7 + 5 \times 2 \;\Big|\; 5 \times 7. \qquad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

This is a very awkward way of writing this

Comments on page 3

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \tag{2.2}$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \tag{2.3}$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \tag{2.4}$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \tag{2.5}$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \tag{2.6}$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \times 2|7 = 3 \times 2 \;\big|\; 3 \times 7 + 5 \times 2 \;\big|\; 5 \times 7. \tag{2.7}$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

Where does this "school method" come from? I was taught to do multiplication by hand this way. 35 x27 —- 245 70 945 (7*5)=35, write down the 5, carry the 3 over, 7*3+3=24, write that down. Then repeat for 2. 2*5=10, write 0, carry 1 over, 2*3+1=7, write that down. Then add the two numbers.

I would definitely say that the above would be reflective of my schooling, however I find myself using this other method (outlined in the text) when I need ot multiply larger numbers, although I wouldn't say I was ever formally taught it. And I never took the time to write it all out. Very interesting where you take this.

I agree that this is how I was taught to do multiplication, but most of us wouldn't do it this way anymore - it would probably be 27*30 + 27*10/2 since that's faster and more intuitive for us.

I agree with this method, it's much faster and how I would approach it, although I did (35x30) - (35x3) which came out to a clean 945.

I agree that this is a bit more intuitive however, i think it takes more effort and concentration to do all of this in ones head and remember the numbers too. Since many human beings are lazy it seems that there is a bit of a trade off of effort vs. time (since you can probably do it in your head faster than on paper).

If you look at what you've written out, I think that's the answer to your question. That is school method except it is written in column form instead of side-ways in regular equation form. The general procedure is to take the units, tens, hundreds (27 = 7+20) and multiply by the first number. When you take each digit and multiply by the first number, you do it by breaking down the first number into units, tens, hundreds (35=30+5). In the end, you multiple out everything keeping track of all powers of ten.

That being said, it might be easier to tie introduce this section using the column multiplication that we all learned.

If I am doing the multiplication on paper, I use the column method as explained above. But if for whatever reason I need to do it in my head, I"ll use the method as explained in the text. (assuming the problem is 2*2 or maybe 3*2 max... after that I can't keep the numbers in my memory easily and will resort to paper or calculator)

Remember from algebra: (3x+5)(2x+7)=... He's applying what your teacher's made you drill in abstract

this is a clever insight, it makes it much easier to understand what he is doing in Eg. 2.4. Old FOIL method.

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \quad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \quad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \quad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \quad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \quad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \ \times \ 2|7 = 3 \times 2 \ \big| \ 3 \times 7 + 5 \times 2 \ \big| \ 5 \times 7. \quad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

school method? did anyone do this in school?

Yes, I did this in school.

I didn't do this in school, but I can see how this method would be taught. I was a bit confused by the name "school method" as well.

it is a horizontal representation of the vertical multiplication that i learned

I didn't realize that. Maybe it would be easier to recognize if it was presented in the vertical form somehow.

I assume the recursion comes in for each n power of 10?

y are we regrouping by the powers of 10? is this in reference to befoe when we estimated large multiplications by separating the power of 10 from the from number?

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots . \quad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \quad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \quad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \quad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \quad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \times 2|7 = 3 \times 2 \;\big|\; 3 \times 7 + 5 \times 2 \;\big|\; 5 \times 7. \quad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

This is a really confusing way to write the math out.

I agree... how is this more useful than standard two-number multiplication where you start with the units digit, and carry the factors of 10, etc.?

Yeah, I would think that the school-method would just be the multiplication method just mentioned? I understand that this method represents recursion but it's also making this multiplication so much more difficult than it is. I though recursion was supposed to simplify our calculations?

I don't think so, I mean it's pretty clear what's happening here in terms of grouping factors and dealing with distribution. I think the target audience of this book is college age or older students / teachers, who shouldn't have that much trouble figuring out the math. Also, the point is kind of to show how cumbersome and totally unnecessary the 'school method' is compared to smart recursion.

I don't think so, I mean it's pretty clear what's happening here in terms of grouping factors and dealing with distribution. I think the target audience of this book is college age or older students / teachers, who shouldn't have that much trouble figuring out the math. Also, the point is kind of to show how cumbersome and totally unnecessary the 'school method' is compared to smart recursion.

I don't think so, I mean it's pretty clear what's happening here in terms of grouping factors and dealing with distribution. I think the target audience of this book is college age or older students / teachers, who shouldn't have that much trouble figuring out the math. Also, the point is kind of to show how cumbersome and totally unnecessary the 'school method' is compared to smart recursion.

I don't think so, I mean it's pretty clear what's happening here in terms of grouping factors and dealing with distribution. I think the target audience of this book is college age or older students / teachers, who shouldn't have that much trouble figuring out the math. Also, the point is kind of to show how cumbersome and totally unnecessary the 'school method' is compared to smart recursion.

I don't think so, I mean it's pretty clear what's happening here in terms of grouping factors and dealing with distribution. I think the target audience of this book is college age or older students / teachers, who shouldn't have that much trouble figuring out the math. Also, the point is kind of to show how cumbersome and totally unnecessary the 'school method' is compared to smart recursion.

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \qquad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \qquad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \qquad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \qquad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \qquad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \times 2|7 = 3 \times 2 \;\big|\; 3 \times 7 + 5 \times 2 \;\big|\; 5 \times 7. \qquad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

While compared to the school method this takes less paper space and the math isnt too hard to figure out, it still comes across as an application of recursion with hard to see benefits. Adding in the following section was very useful though.

could you at least add a few more spaces between the terms being added? i felt myself goign back several times and trying to find the term. it made reading it take way longer than necessary.

Or, as suggested in another comment, I will add parentheses to make the grouping clear.

This totally glitched out and over-posted a bunch of times. Weird.

Maybe add a feature to delete double posts?

is that a reliable definition of an abstraction?

I think so... its reusable and keeps only important details.

This is similar to what he was talking about in class with the tree programming language

Actually, the programming the tree thing makes me wonder – is object oriented coding all about abstraction? You abstract things into objects, classes, etc...?

I don't think that this is meant to be a definition of abstraction, it is only pointing out that commonly convenient notation's are also abstractions.

A notation is not the only kind of abstraction, but it's one of the most useful and easy to recognize (and common).

Think of musical notation, e.g. for piano or guitar music. It is such a good notation that we hardly even realize that it is a notation.

i dont understand this type of notation

10y+x, I think?

agreed!

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \quad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \quad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \quad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \quad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \quad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x+y$ and $z|y|x$ represent $100z+10y+x$. Then the school method runs as follows:

$$3|5 \times 2|7 = 3 \times 2 \;\big|\; 3 \times 7 + 5 \times 2 \;\big|\; 5 \times 7. \quad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

I believe this should be:

y|x represents 10y+x.

As it is different from the next example and doesn't make sense in equation 2.7

I agree I think this might have been a typo, otherwise 3|5 would yield 53 and not 35

Perhaps a better way to present this abstraction would be to present the calculation graphically, similar to how the alternative "school method" mentioned in the comments above is organized? (I mean, by lining up all the terms that are x1, and all the x10 terms, and all the x100 terms, etc, then simply adding them rather than multiplying them?) It would avoid the cluttering of the | syntax here.

how would this work for fractions?

For decimal expansions you could maybe do: 3|1|.4|1|5 for 31.415? I'm not sure why it would need to work for fractions..

i am honestly only glossing over this page and all the arithmetic manipulations and labeling it as an inefficient method. im not sure if that is the point of this, but this is how it is coming across to me.

This is rather confusing when it's embedded in the text. If it's going to be so crucial for the next part, try separating it out of the text.

How do you know where the |'s are used to separate the numbers you multiply?

How did this come from that description? I would have expected (10*3+5)*(2*10+7) and so on

Problem 2.1 Summing a series using abstraction
Use abstraction to find the sum of the infinite series

$$1 + r + r^2 + r^3 + \cdots. \quad (2.2)$$

2.2.2 Computational Recursion

The second example of recursion is an algorithm to multiply many-digit numbers much more rapidly than is possible with the standard school method. The school method is sufficient for humans, for we rarely multiply large numbers by hand. However, computers are often called upon to multiply gigantic numbers, whether in computing π to billions of digits or in public-key cryptography. I'll introduce the new method by contrasting it with the school method on the example of 35×27.

In the school method, the product is written as

$$35 \times 27 = (3 \times 10 + 5) \times (2 \times 10 + 7). \quad (2.3)$$

The product expands into four terms:

$$(3 \times 10) \times (2 \times 10) + (3 \times 10) \times 7 + 5 \times (2 \times 10) + 5 \times 7. \quad (2.4)$$

Regrouping the terms by the powers of 10 gives

$$3 \times 2 \times 100 + (3 \times 7 + 5 \times 2) \times 10 + 5 \times 7. \quad (2.5)$$

Then you remember the four one-digit multiplications 3×2, 3×7, 5×2, and 5×7, finding that

$$35 \times 27 = 6 \times 100 + 31 \times 10 + 35 = 945. \quad (2.6)$$

Unfortunately, the preceding description is cluttered with powers of 10 obscuring the underlying pattern. Therefore, define a convenient notation (an abstraction!): Let $y|x$ represent $10x + y$ and $z|y|x$ represent $100z + 10y + x$. Then the school method runs as follows:

$$3|5 \times 2|7 = 3 \times 2 \mid 3 \times 7 + 5 \times 2 \mid 5 \times 7. \quad (2.7)$$

This notation shows how school multiplication replaces a two-digit multiplication with four one-digit multiplications. It would recursively replace

It took me a long time to figure out what this was saying. As written, I read (3x2)x10 + 3x7 + (5x2)x10 + 5x7. I think it'd be much more clear if you used parentheses, i.e. (3x2)|(3x7+5x2)|(5x7). The root of this problem is you never specified where the | feel with regards to order of operations. Also, y|x is meant in discrete mathematics (at least as taught in 6.042) that x divides y (y=ax), so this might be confusing.

I agree. I think a few sentences explaining the exact method of getting to that expression would be helpful. Especially during its first appearance in this chapter. For example the expression in the middle bracket wasn't clearly apparent.

Agreed, I had to read over this section several times in order to understand what was meant by |. Even just a quick explanation of this would and how you came up with the notation would be very helpful.

Need to see this in class. I don't get it.

Need to see this in class. I don't get it.

I'm not sure if I get it, which is also why I would like to see it in class.

I had never thought of multiplication in this manner; this is really useful, and I think after some practice with it, it will be a powerful tool.

I agree - parentheses would've helped emphasize that you are doing (x)|(y)|(z)

I think I'm still a little confused about what is going on here... I can't really see how this method is supposed easier or better than other methods.

It also took me a little while to see what this was saying. But something I noticed is that the expansion of the multiplication is kind of like the FOIL method taught for expanding multiplication of factors.

why is this helpful? it seems like the same thing as above

I need some more time to learn to do this.

a four-digit multiplication with four two-digit multiplications. For example, using a modified | notation where y|x means 100y + x, the product 3247 × 1798 becomes

$$32|47 \times 17|98 = 32 \times 17 \;\Big|\; 32 \times 98 + 47 \times 17 \;\Big|\; 47 \times 98. \qquad (2.8)$$

Each two-digit multiplication (of which there are four) would in turn become four one-digit multiplications. For example (and using the normal y|x = 10y + x notation),

$$3|2 \times 1|7 = 3 \times 2 \;\Big|\; 3 \times 7 + 2 \times 1 \;\Big|\; 2 \times 7. \qquad (2.9)$$

Thus, a four-digit multiplication becomes 16 one-digit multiplications.

Continuing the pattern, an eight-digit multiplication becomes four four-digit multiplications or, in the end, 64 one-digit multiplications. In general, an n-digit multiplication requires n^2 one-digit multiplications. This recursive algorithm seems so natural, perhaps because we learned it so long ago, that improvements are hard to imagine.

Surprisingly, a slight change in the method significantly improves it. The key is to retain the core idea of recursion but to improve the method of decomposition. Here is the improvement:

$$a_1|a_0 \times b_1|b_0 = a_1b_1 \;\Big|\; (a_1 + a_0)(b_1 + b_0) - a_1b_1 - a_0b_0 \;\Big|\; a_0b_0.$$

Before analyzing the improvement, let's check that it is not nonsense by retrying the 35 × 27 example.

$$3|5 \times 2|7 = 3 \times 2 \;\Big|\; (3+5)(2+7) - 3 \times 2 - 5 \times 7 \;\Big|\; 5 \times 7.$$

Doing the five one-digit multiplications gives

$$3|5 \times 2|7 = 6\,|\,31\,|\,35 = 6 \times 100 + 31 \times 10 + 35 = 945, \qquad (2.10)$$

just as it should.

At first glance, the method seems like a retrograde step because it requires five multiplications whereas the school method requires only four. However, the magic of the new method is that two multiplications are redundant: a_1b_1 and a_0b_0 are each computed twice. Therefore, the new method requires only three multiplications. The small change from four to three multiplications, when used recursively, makes the new method significantly faster: An n-digit multiplication requires roughly $n^{1.58}$ one-digit

Comments on page 4

I might recommend using a different notation here such that

y||x means 100y + x

Agreed. It would make things more clear so that instead of trying to figure out whether | means 10y+x or 100y+x, we can just focus on the more relevant parts of the problem.

Yeah this is a little confusing.

first thing that entered my mind I was thinking "given" like y given x, since we were just talking about probabilities

I agree. My first reaction was that this must have been a typo, and your change to the notation would make it much clearer that it's intentional.

Looks very cluttered and hard to read

As soon as you understand the method it's really not that cluttered, I don't think there's a better way to illustrate the multiplication.

I feel like this makes the problem even more complicated

I think this looks fine. It flows well. Maybe you could make the bars bolder.

here it is corrected

You could possibly use arrows to make things clearer? (how the multiplying is done etc.

I think that's a typo. Should be 3 x 1|3 x 7...

a four-digit multiplication with four two-digit multiplications. For example, using a modified | notation where $y|x$ means $100y + x$, the product 3247×1798 becomes

$$32|47 \times 17|98 = 32 \times 17 \mid 32 \times 98 + 47 \times 17 \mid 47 \times 98. \quad (2.8)$$

Each two-digit multiplication (of which there are four) would in turn become four one-digit multiplications. For example (and using the normal $y|x = 10y + x$ notation),

$$3|2 \times 1|7 = 3 \times 2 \mid 3 \times 7 + 2 \times 1 \mid 2 \times 7. \quad (2.9)$$

Thus, a four-digit multiplication becomes 16 one-digit multiplications.

Continuing the pattern, an eight-digit multiplication becomes four four-digit multiplications or, in the end, 64 one-digit multiplications. In general, an n-digit multiplication requires n^2 one-digit multiplications. This recursive algorithm seems so natural, perhaps because we learned it so long ago, that improvements are hard to imagine.

Surprisingly, a slight change in the method significantly improves it. The key is to retain the core idea of recursion but to improve the method of decomposition. Here is the improvement:

$$a_1|a_0 \times b_1|b_0 = a_1b_1 \mid (a_1 + a_0)(b_1 + b_0) - a_1b_1 - a_0b_0 \mid a_0b_0.$$

Before analyzing the improvement, let's check that it is not nonsense by retrying the 35×27 example.

$$3|5 \times 2|7 = 3 \times 2 \mid (3 + 5)(2 + 7) - 3 \times 2 - 5 \times 7 \mid 5 \times 7.$$

Doing the five one-digit multiplications gives

$$3|5 \times 2|7 = 6 \mid 31 \mid 35 = 6 \times 100 + 31 \times 10 + 35 = 945, \quad (2.10)$$

just as it should.

At first glance, the method seems like a retrograde step because it requires five multiplications whereas the school method requires only four. However, the magic of the new method is that two multiplications are redundant: a_1b_1 and a_0b_0 are each computed twice. Therefore, the new method requires only three multiplications. The small change from four to three multiplications, when used recursively, makes the new method significantly faster: An n-digit multiplication requires roughly $n^{1.58}$ one-digit

Does this method make multiplication easy to do in your head? It seems like 16 1-digit multiplications would be easy to lose track of, since there's a limit to how many digits people can keep in their temporary memory at a time.

I agree that breaking the problem into so many pieces may make it hard to keep track of the entire problem. Though I haven't quite wrapped my head around this method yet, would it be possible to work through the problem as you're breaking it down. So instead of keeping all 16 multiplications in your head, you work through a part at a time (keeping track of what you still have to break down before the next step)?

I would personally break this down further and do it in parts. I agree that it does seem a bit difficult to keep track of everything.

I would personally break this down further and do it in parts. I agree that it does seem a bit difficult to keep track of everything.

I think the point is to create problems which we can solve in our head more easily rather then make it easier to keep track of. Clearly, breaking the numbers down to smaller numbers makes more solutions to keep track of but it's a good balance that we're after.

I found this very very interesting. One of the things I have most enjoyed reading about thus far.

I feel the same way. Wouldn't it be easier to just use divide and conquer?

Pretty neat. But doesn't this make room for a lot of mistakes, going against our idea of intelligent redundancy?

is this faster than the school method?

some sort of comparison?

i never learned multiplication looking like this... doesn't seem as natural, but it definitely makes sense

a four-digit multiplication with four two-digit multiplications. For example, using a modified | notation where $y|x$ means $100y + x$, the product 3247×1798 becomes

$$32|47 \times 17|98 = 32 \times 17 \mid 32 \times 98 + 47 \times 17 \mid 47 \times 98. \quad (2.8)$$

Each two-digit multiplication (of which there are four) would in turn become four one-digit multiplications. For example (and using the normal $y|x = 10y + x$ notation),

$$3|2 \times 1|7 = 3 \times 2 \mid 3 \times 7 + 2 \times 1 \mid 2 \times 7. \quad (2.9)$$

Thus, a four-digit multiplication becomes 16 one-digit multiplications.

Continuing the pattern, an eight-digit multiplication becomes four four-digit multiplications or, in the end, 64 one-digit multiplications. In general, an n-digit multiplication requires n^2 one-digit multiplications. This recursive algorithm seems so natural, perhaps because we learned it so long ago, that improvements are hard to imagine.

Surprisingly, a slight change in the method significantly improves it. The key is to retain the core idea of recursion but to improve the method of decomposition. Here is the improvement:

$$a_1|a_0 \times b_1|b_0 = a_1b_1 \mid (a_1 + a_0)(b_1 + b_0) - a_1b_1 - a_0b_0 \mid a_0b_0.$$

Before analyzing the improvement, let's check that it is not nonsense by retrying the 35×27 example.

$$3|5 \times 2|7 = 3 \times 2 \mid (3 + 5)(2 + 7) - 3 \times 2 - 5 \times 7 \mid 5 \times 7.$$

Doing the five one-digit multiplications gives

$$3|5 \times 2|7 = 6 \,|\, 31 \,|\, 35 = 6 \times 100 + 31 \times 10 + 35 = 945, \quad (2.10)$$

just as it should.

At first glance, the method seems like a retrograde step because it requires five multiplications whereas the school method requires only four. However, the magic of the new method is that two multiplications are redundant: a_1b_1 and a_0b_0 are each computed twice. Therefore, the new method requires only three multiplications. The small change from four to three multiplications, when used recursively, makes the new method significantly faster: An n-digit multiplication requires roughly $n^{1.58}$ one-digit

I don't really see this algorithm as "natural" though, it is more complicated to me than just multiplying via one number on top of the other and carrying the 10s, etc.

yeah I agree, although this method makes sense, I was definitely not taught to do multiplication this way at school so it doesn't feel "natural" as stated in the paragraph.

Well, I think it feels unnatural because it sounds like most of us never actually learned multiplication this way (this is the first time I've seen this)... so it feels like there is some sort of gap here based on the way people were taught multiplication. I can see where this is going in terms of recursion, but it certainly isn't the way I'd multiply by hand.

Maybe if we were taught multiplication by this recursion method, it would seem more natural. If nothing else, introducing the concept of recursion earlier in schooling would allow people to understand this concept better

How does this improve the method?

visual diagrams are actually a lot more helpful for me rather than writing out the recursion steps

Why would you break it down into this step? It seems more complicated?

I think these examples are interesting and do illustrate the recursion and abstraction principles, but I am still having a hard time relating them to something a person would do rather than a computer. Maybe the section should include an example that is practical for problem solving by the human brain.

i agree...it's sorta interesting, but i don't think i would ever do this

I can see how 2.9 is very useful but the improvement seems a lot more complicated to remember for a four or six digit multiplication.

How many orders of computation are being saved?

This is a really cool method.

When doing quick approximations, could we speed up the process by adding numbers to make "few"?

i like the logic used here. not sure if it is necessary though, as it is assumed that since it is in the book the new method should work by most students.

a four-digit multiplication with four two-digit multiplications. For example, using a modified | notation where $y|x$ means $100y + x$, the product 3247×1798 becomes

$$32|47 \times 17|98 = 32 \times 17 \mid 32 \times 98 + 47 \times 17 \mid 47 \times 98. \quad (2.8)$$

Each two-digit multiplication (of which there are four) would in turn become four one-digit multiplications. For example (and using the normal $y|x = 10y + x$ notation),

$$3|2 \times 1|7 = 3 \times 2 \mid 3 \times 7 + 2 \times 1 \mid 2 \times 7. \quad (2.9)$$

Thus, a four-digit multiplication becomes 16 one-digit multiplications.

Continuing the pattern, an eight-digit multiplication becomes four four-digit multiplications or, in the end, 64 one-digit multiplications. In general, an n-digit multiplication requires n^2 one-digit multiplications. This recursive algorithm seems so natural, perhaps because we learned it so long ago, that improvements are hard to imagine.

Surprisingly, a slight change in the method significantly improves it. The key is to retain the core idea of recursion but to improve the method of decomposition. Here is the improvement:

$$a_1|a_0 \times b_1|b_0 = a_1b_1 \mid (a_1 + a_0)(b_1 + b_0) - a_1b_1 - a_0b_0 \mid a_0b_0.$$

Before analyzing the improvement, let's check that it is not nonsense by retrying the 35×27 example.

$$3|5 \times 2|7 = 3 \times 2 \mid (3 + 5)(2 + 7) - 3 \times 2 - 5 \times 7 \mid 5 \times 7.$$

Doing the five one-digit multiplications gives

$$3|5 \times 2|7 = 6 \,|\, 31 \,|\, 35 = 6 \times 100 + 31 \times 10 + 35 = 945, \quad (2.10)$$

just as it should.

At first glance, the method seems like a retrograde step because it requires five multiplications whereas the school method requires only four. However, the magic of the new method is that two multiplications are redundant: a_1b_1 and a_0b_0 are each computed twice. Therefore, the new method requires only three multiplications. The small change from four to three multiplications, when used recursively, makes the new method significantly faster: An n-digit multiplication requires roughly $n^{1.58}$ one-digit

Wow this is crazy how it works, but I still don't totally understand exactly how it works

I'm also a little lost. Is there some way to convey this graphically? Which numbers are being multiplied where and why.

I feel like this relates a lot to the chess game. Yes this is great for a computer because they dont have trouble remembering but for humans it doesnt seem practical. I need more convincing.

It is definitely not practical for a human! The method is used here to illustrate how abstraction, of which recursion is a special case, leads to understanding a very sly algorithm. The goal of the course is to understand lots of natural systems (e.g. blue skies) and to learn tools that help in designing and building person-made systems (e.g. bridges, large software systems).

Wow. That is really impressive. I never would've though of that

Woah, this is also amazing. These alternative methods of using abstraction are very interesting.

I was wondering where we were saving time.

I feel like this is true for programming. The computer will notice that there were a few unnecessary calculations. But for humans to realize that would take more time. Recursion to me doesn't seem like a practical method for humans.

This is awesome. Why don't they teach this to us at school?

Ok this makes way more sense as to why it is "clever redundancy!"

I agree, I wish I had learned this a long time ago. I would assume that it is even more practical for humans because of the need to eliminate calculations while working with mental math.

a four-digit multiplication with four two-digit multiplications. For example, using a modified | notation where $y|x$ means $100y + x$, the product 3247×1798 becomes

$$32|47 \times 17|98 = 32 \times 17 \mid 32 \times 98 + 47 \times 17 \mid 47 \times 98. \quad (2.8)$$

Each two-digit multiplication (of which there are four) would in turn become four one-digit multiplications. For example (and using the normal $y|x = 10y + x$ notation),

$$3|2 \times 1|7 = 3 \times 2 \mid 3 \times 7 + 2 \times 1 \mid 2 \times 7. \quad (2.9)$$

Thus, a four-digit multiplication becomes 16 one-digit multiplications.

Continuing the pattern, an eight-digit multiplication becomes four four-digit multiplications or, in the end, 64 one-digit multiplications. In general, an n-digit multiplication requires n^2 one-digit multiplications. This recursive algorithm seems so natural, perhaps because we learned it so long ago, that improvements are hard to imagine.

Surprisingly, a slight change in the method significantly improves it. The key is to retain the core idea of recursion but to improve the method of decomposition. Here is the improvement:

$$a_1|a_0 \times b_1|b_0 = a_1b_1 \mid (a_1 + a_0)(b_1 + b_0) - a_1b_1 - a_0b_0 \mid a_0b_0.$$

Before analyzing the improvement, let's check that it is not nonsense by retrying the 35×27 example.

$$3|5 \times 2|7 = 3 \times 2 \mid (3+5)(2+7) - 3 \times 2 - 5 \times 7 \mid 5 \times 7.$$

Doing the five one-digit multiplications gives

$$3|5 \times 2|7 = 6 \mid 31 \mid 35 = 6 \times 100 + 31 \times 10 + 35 = 945, \quad (2.10)$$

just as it should.

At first glance, the method seems like a retrograde step because it requires five multiplications whereas the school method requires only four. However, the magic of the new method is that two multiplications are redundant: a_1b_1 and a_0b_0 are each computed twice. Therefore, the new method requires only three multiplications. The small change from four to three multiplications, when used recursively, makes the new method significantly faster: An n-digit multiplication requires roughly $n^{1.58}$ one-digit

I don't understand how this change would really make a computer faster. The way I see it, a computer would have a database with all one digit multiplications and their results. My sense of the reading is that the time-saving comes from the fact that while there are the same number of 1x1 multiplications, some of them are the duplicates. I don't understand how it would take less time for it to "compute" this number with the duplicates than the case without them, wouldn't it still need to replace the same number of 1x1 calculations with answers?

This is actually really clever. At first it seems very complicated, but if you can follow their recursion, it actually makes it a lot faster

Yeah, once I understood that a1b1 and a0b0 are the same, it made sense.

where does the 1.58 come from? is that the log2(3) as mentioned below?

Why is it log2(3) though?

multiplications (Problem 2.2). In contrast, the school algorithm requires n^2 one-digit multiplications. The small decrease in the exponent from 2 to 1.58 has a large effect when n is large. For example, when multiplying billion-digit numbers, the ratio of n^2 to $n^{\log_2 3}$ is roughly 5000.

Why would anyone multiply billion-digit numbers? One answer is to compute π to a billion digits. Computing π to a huge number of digits, and comparing the result with the calculations of other supercomputers, is the standard way to verify the numerical hardware in a new supercomputer.

The new algorithm is known as the Karatsuba algorithm after its inventor [15]. But even it is too slow for gigantic numbers. For large enough n, an algorithm using fast Fourier transforms is even faster than the Karatsuba algorithm. The so-called Schönhage–Strassen algorithm [27] requires a time proportional to $n \log n \log\log n$. High-quality libraries for large-number multiplication recursively use a combination of regular multiplication, Karatsuba, and Schönhage–Strassen, selecting the algorithm according to the number of digits.

Problem 2.2 Running time of the Karatsuba algorithm
Show that the Karatsuba multiplication method requires $n^{\log_2 3} \approx n^{1.58}$ one-digit multiplications.

2.3 Low-pass filters

2.3.1 RC circuits

2.3.2 Light-bulb flicker

2.3.3 Temperature fluctuations

2.4 Summary and further problems

The diagram for the hiker has two names: a phase-space diagram or a spacetime diagram. Both types are useful in science and engineering. Spacetime diagrams, used in Einstein's theory of relativity, are the subject of the wonderful textbook [30]. They are the essential ingredient in

Very interesting.. Who originally thought of this? Why do we learn the other way?

Yea school way is much less efficient than computational recursion

but it's less complicated to learn! we don't really learn the school method in the same way that it's explained here...we do it vertically. This physical setup is much easier to remember [and learn than] an equation

is there an example of it being used for higher order multiplication?

is there an example of it being used for higher order multiplication?

How does multiplying billion digits allow one to compute pi? How do people calculate pi anyways?

http://en.wikipedia.org/wiki/Numerical_approximations_of_%CF%80

I'm really glad you asked that because that was a really interesting article... I think the question you asked has a million different answers and has been an important questoin for a long time

I like that you explain a concept then tell us why this is important at all.

In the smaller form, this feels useful for humans to possibly use as a way to do 2 digit (maybe 4 digit) multiplication. Even though I think it is not useful for humans (we have calculators) past that, it's nice to now know how this type of thing works. Very informational; I think this section was helpful.

very interesting fact, i never knew why people wanted to compute pi to a more digits than 10

multiplications (Problem 2.2). In contrast, the school algorithm requires n^2 one-digit multiplications. The small decrease in the exponent from 2 to 1.58 has a large effect when n is large. For example, when multiplying billion-digit numbers, the ratio of n^2 to $n^{\log_2 3}$ is roughly 5000.

Why would anyone multiply billion-digit numbers? One answer is to compute π to a billion digits. Computing π to a huge number of digits, and comparing the result with the calculations of other supercomputers, is the standard way to verify the numerical hardware in a new supercomputer.

The new algorithm is known as the Karatsuba algorithm after its inventor [15]. But even it is too slow for gigantic numbers. For large enough n, an algorithm using fast Fourier transforms is even faster than the Karatsuba algorithm. The so-called Schönhage–Strassen algorithm [27] requires a time proportional to $n \log n \log \log n$. High-quality libraries for large-number multiplication recursively use a combination of regular multiplication, Karatsuba, and Schönhage–Strassen, selecting the algorithm according to the number of digits.

Problem 2.2 Running time of the Karatsuba algorithm
Show that the Karatsuba multiplication method requires $n^{\log_2 3} \approx n^{1.58}$ one-digit multiplications.

2.3 Low-pass filters

2.3.1 RC circuits

2.3.2 Light-bulb flicker

2.3.3 Temperature fluctuations

2.4 Summary and further problems

The diagram for the hiker has two names: a phase-space diagram or a spacetime diagram. Both types are useful in science and engineering. Spacetime diagrams, used in Einstein's theory of relativity, are the subject of the wonderful textbook [30]. They are the essential ingredient in

Random, but interesting. I like these little sidenotes, they help me remember what I read.

That's true, it is useful how the random application examples help differentiate topics.

I agree I would encourage including these random interesting facts because the random/interesting nature makes them easy to remember and thus it is easier to remember the material associated with them.

I'm going to have to fourth this. It's a lot of these side notes, ans short purposes for the readings that keeps me engaged and saying "I wonder what this is used for? Ohhh!"

Yeah, these are really cool to read over.

That answers my question..

It's not that clear that the Karatsuba algorithm refers to what we just did above. I had to Wikipedia that...

I don't agree, I think it is clear that the "new algorithm" refers to the algorithm that was just explained in the previous paragraphs.

I'm kind of confused as to how this applies to recursion. Another example after explaining it would be helpful.

I somewhat agree... I thought it was a really interesting side note as someone who is pretty technically competent in these types of things. I could see readers getting pretty lost though.

I was definitely confused a little bit here but I do see how it relates to abstraction/recusrion

multiplications (Problem 2.2). In contrast, the school algorithm requires n^2 one-digit multiplications. The small decrease in the exponent from 2 to 1.58 has a large effect when n is large. For example, when multiplying billion-digit numbers, the ratio of n^2 to $n^{\log_2 3}$ is roughly 5000.

Why would anyone multiply billion-digit numbers? One answer is to compute π to a billion digits. Computing π to a huge number of digits, and comparing the result with the calculations of other supercomputers, is the standard way to verify the numerical hardware in a new supercomputer.

The new algorithm is known as the Karatsuba algorithm after its inventor [15]. But even it is too slow for gigantic numbers. For large enough n, an algorithm using fast Fourier transforms is even faster than the Karatsuba algorithm. The so-called Schönhage–Strassen algorithm [27] requires a time proportional to $n \log n \log \log n$. High-quality libraries for large-number multiplication recursively use a combination of regular multiplication, Karatsuba, and Schönhage–Strassen, selecting the algorithm according to the number of digits.

Problem 2.2 Running time of the Karatsuba algorithm
Show that the Karatsuba multiplication method requires $n^{\log_2 3} \approx n^{1.58}$ one-digit multiplications.

2.3 Low-pass filters

2.3.1 RC circuits

2.3.2 Light-bulb flicker

2.3.3 Temperature fluctuations

2.4 Summary and further problems

The diagram for the hiker has two names: a phase-space diagram or a spacetime diagram. Both types are useful in science and engineering. Spacetime diagrams, used in Einstein's theory of relativity, are the subject of the wonderful textbook [30]. They are the essential ingredient in

Are we going to actually see this later? Otherwise it seems slightly unnecessary.

Agreed - I don't think algorithm names should be thrown around unless they'll come back later.

I think a one liner name doesn't hurt to be put in here, it gives you a bit of insight into this world of algorithms an you can research it if you want on your own.

but it takes away from the message here and distracts the reader (clearly)

I think it's a nice addition to reference additional algorithms at the end of a section. It gives more background on the topic and if someone is interested in learning about more complex/efficient algorithms they can go look them up or at least know about their existence.

Could someone please explain the background on algorithm speeds and the meaning of "n logn loglogn"?

yea i'd like one too...unless it's ridiculously long and complicated, then i guess i'd rather know it's not worth my time

It would be a little difficult to explain under these circumstances, but take a look at the Wikipedia article under "big O Notation" hopefully it would be a bit cleaer from there.

This is unclear where the parentheses are.

n (log n)(log (log n))

The multiple terms probably derive from multiple loops within the algorithm and/or program.

How?? if this SS algorithm is so fast, why isn't the only one used, and maybe some info on how SS works?

multiplications (Problem 2.2). In contrast, the school algorithm requires n^2 one-digit multiplications. The small decrease in the exponent from 2 to 1.58 has a large effect when n is large. For example, when multiplying billion-digit numbers, the ratio of n^2 to $n^{\log_2 3}$ is roughly 5000.

Why would anyone multiply billion-digit numbers? One answer is to compute π to a billion digits. Computing π to a huge number of digits, and comparing the result with the calculations of other supercomputers, is the standard way to verify the numerical hardware in a new supercomputer.

The new algorithm is known as the Karatsuba algorithm after its inventor [15]. But even it is too slow for gigantic numbers. For large enough n, an algorithm using fast Fourier transforms is even faster than the Karatsuba algorithm. The so-called Schönhage–Strassen algorithm [27] requires a time proportional to $n \log n \log\log n$. High-quality libraries for large-number multiplication recursively use a combination of regular multiplication, Karatsuba, and Schönhage–Strassen, selecting the algorithm according to the number of digits.

Problem 2.2 Running time of the Karatsuba algorithm
Show that the Karatsuba multiplication method requires $n^{\log_2 3} \approx n^{1.58}$ one-digit multiplications.

2.3 Low-pass filters

2.3.1 RC circuits

2.3.2 Light-bulb flicker

2.3.3 Temperature fluctuations

2.4 Summary and further problems

The diagram for the hiker has two names: a phase-space diagram or a spacetime diagram. Both types are useful in science and engineering. Spacetime diagrams, used in Einstein's theory of relativity, are the subject of the wonderful textbook [30]. They are the essential ingredient in

Is there a reference to where we can read about this in more detail? it sounds very interesting.

one of my favorite things about this class is that in addition to just "approximation" we get to learn all sorts of interesting things... such as why lectures are the way they are, how CDs work, and how pi can be computed faster :)

I find it very interesting that various programs will look at a problem and decide the best way of computing the answer based on the number of digits. I didn't realize programs are so dynamic.

The GNU Multiple Precision library (http://gmplib.org/) is a very high-quality library that does that. The page also has a link to a program (using the library) that computes pi to one billion digits.

How do I know that it chooses the algorithm by the size of the numbers? I think I remember reading about that in the documentation, but I also experimented with it. Python has an interface (a "binding") to the library, and plotted the running times versus number size, and you can see the breakpoints in the graph as the algorithm changes.

Maybe we can go over Strassen's Algorithm for matrix multiplication? It would be nice to find an abstraction there.

can you please go over this method in class? I am having a really hard time reading about it and understanding what is going on

While it should be clear what n is based on the above discussion, it would still be nice to explicitly say that we are calculating O(n) for a n-digit times a n-digit number.

agreed. does this have to do with abstraction directly as well? or does it have to do with the algorithm that used abstraction a while ago?

I think the examples given above could be stronger by adding a closing paragraph about how it directly relates to abstraction. I feel like the term abstraction was somewhat avoided in this section.

Overall an informative and interesting chapter. It is generally well written (though adding parentheses to the | notation would improve it a lot).

multiplications (Problem 2.2). In contrast, the school algorithm requires n^2 one-digit multiplications. The small decrease in the exponent from 2 to 1.58 has a large effect when n is large. For example, when multiplying billion-digit numbers, the ratio of n^2 to $n^{\log_2 3}$ is roughly 5000.

Why would anyone multiply billion-digit numbers? One answer is to compute π to a billion digits. Computing π to a huge number of digits, and comparing the result with the calculations of other supercomputers, is the standard way to verify the numerical hardware in a new supercomputer.

The new algorithm is known as the Karatsuba algorithm after its inventor [15]. But even it is too slow for gigantic numbers. For large enough n, an algorithm using fast Fourier transforms is even faster than the Karatsuba algorithm. The so-called Schönhage–Strassen algorithm [27] requires a time proportional to $n \log n \log \log n$. High-quality libraries for large-number multiplication recursively use a combination of regular multiplication, Karatsuba, and Schönhage–Strassen, selecting the algorithm according to the number of digits.

Problem 2.2 Running time of the Karatsuba algorithm
Show that the Karatsuba multiplication method requires $n^{\log_2 3} \approx n^{1.58}$ one-digit multiplications.

2.3 Low-pass filters

2.3.1 RC circuits

2.3.2 Light-bulb flicker

2.3.3 Temperature fluctuations

2.4 Summary and further problems

The diagram for the hiker has two names: a phase-space diagram or a spacetime diagram. Both types are useful in science and engineering. Spacetime diagrams, used in Einstein's theory of relativity, are the subject of the wonderful textbook [30]. They are the essential ingredient in

Is section 2.3 not written yet? lol. I guess this would emphasize to everyone that this book is a process in the making and we should try to set aside any frustrations by focusing more on constructive comments.

I was tempted to put a page break in before Section 2.3, to avoid showing my hand. But I resisted. Indeed, in some parts of the course, I have a larger margin of safety than in others.

This unit on abstraction is the one where I have the least margin of safety. I've been thinking about it and trying different versions for a few years, and this year has been the most coherent so far (but I leave to you to decide whether, on an absolute scale, it is actually coherent).

what happened to these guys?

lol don't complain

The posted assignment was to read section 2.2 only so he'll probably add these in before Tuesday

I never thought of light-bulb flicker as a low pass filter. Live and learn, I guess

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

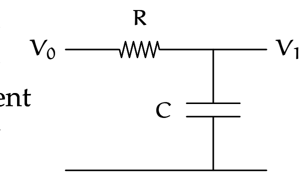

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

are input/output signals always going to be so readily recognizable?

I think a tree would be useful here to show all the factors that can affect resistance to daily temp variations

I agree, by the time you reach the end of this section you realize that these two situations parallel nicely.

Now we can do the last 2 problems of the pset! woohoo!

I enjoyed this chapter because I understood all of it after taking 2.005 and 6.002, which is cool!

It seems like this chapter combines abstraction with divide and conquer. It seems like a really practical application based on the example.

WHy is this a transfer function? Is B a Laplace transform?

Thanks for incorporating course 2 material (ie. 2.005) and course 6 material. I think it makes it easier for both majors in the class to understand.

Are these approximations taken from your book of constants? Or like your usual general estimates?

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

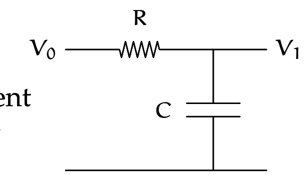

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \qquad (2.11)$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \qquad (2.12)$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

Read Section 2.3 (memo due Tuesday at 10pm). Lots to discuss about it in class, I am sure!

In a section like this, I didn't really see that many points to make. I think most of the other students just made comments about good sentences or summaries. I also thought the section was well done, so I don't have much to say.

Something telling what ontological bonds you are referring to would help here as well as something giving background will help.

Ontology deals with questions concerning what entities exist, and how such entities can be grouped, related within a hierarchy, and subdivided according to similarities and differences. I agree that a definition would be helpful, but I think the definition is implied in the phrase the follows afterwards.

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

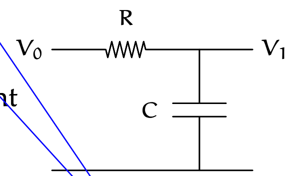

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

Which ones exactly? This sentence made me stop reading and try and figure out which bonds you are referring to..circuits in general? Abstraction?

yeah, using more common vocabulary would make this better. I spent more time trying to figure out what "ontological" meant than necessary.

Agreed, I'm a big fan of new and interesting words, but some like this one just seem misplaced.

I'm in complete agreement with these comments - I did a reading of the sentence, stopped, then went back and read it again. Using that word unnecessarily disrupts the flow.

Is this really that important? I think people are starting to worry too much about grammar and vocabulary

I think it is important because it means that we are having a harder time reading the book. Not to mention, one of the points of us reading is this is to 'provide another set of eyes' and proof read as well...see the how to do reading memos page.

I think this is actually a good place to use the word, the words in between the dashes are another way to say what he means by snipping the ontological bonds. Although I've never heard the word ontological before, I had a good idea of what it meant from the context, I think this is well done.

I think it'd be better to include the words in quotes first, then use the word ontological. It allows the reader to already know the meaning of the word without having to stop focusing on the material.

yeah, the last response hit the spot, reversing the order of the sentence may create less pause and confusion.

those? this doesn't make sense.

Referring perhaps to the bonds of thinking about the study of circuits

definitely going to have to look that one up

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

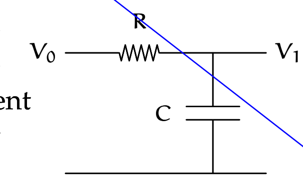

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \quad (2.11)$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \quad (2.12)$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

I understand what you are saying, but I feel that its rather confusingly worded and took me twice reading it to understand

I agree with this completely. I feel like you could condense this whole paragraph down into a much simpler sentence, something to the effect of: "The next example comes from the study of circuits, however we will see that the core idea (the abstraction) will be useful in understanding diverse natural phenomena"

I really like the rephrased comment and greatly agree that the sentence in the reading is overtly confusing. Ontological makes the reading sound sophisticated but is very different from the other vocabulary used in the reading. I personally think it detracts more than it adds

I agree. I got confused from the use of the word those, not really sure what it's referring to. Overall, the paragraph brings up an important point about abstraction.

I think it's the combination of the hyphens, quotes, and parentheses, which really break up the flow of the sentence.

From whence this quote–the dictionary? The quotation marks seem unnecessary unless it's a critical citation.

I agree with the other comments about the confusing sentence, but I also think that after reading this entire paragraph I am almost entirely lost on what you are actually trying to say, besides the idea that you introducing the examples you are about to discuss.

I agree. I think this paragraph does not do a very good job of preparing the reader for what is coming next - it feels wordy in a convoluted way, but doesn't really set up the idea of a low-pass filter, as is the section title.

I disagree. He is explaining that he will use the idea that came from low pass filters and apply it to other things like temperature fluctuations. It motivates the discussion of low pass filters by telling you how it is an example of abstraction.

comma

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

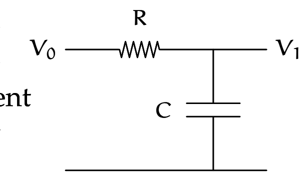

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

After reading this section, I just have to say that I'm glad I'm 6-3, and not 6-1 or 6-2. I don't know if it's the material or the way it was presented, but I pretty confused throughout this whole section.

now think how the course 2 people feel...

Now think how the non course 2 and 6 people feel...

Now think how the non course 2 and 6 people feel...

I'm getting a little annoyed by this I'm not course 6, I dont get it thing...not everyone in course 6 knows how to code [i'm just as much at a loss for that unix stuff as a below average person in an athena cluster...more than joe-shmo on the street, but not enough to get 98% of what he's said in this book.]

my point, please stop insinuating that your knowledge gaps are totally the fault of your major. It's really about your interests...half the course 2 people I know are better at unix than I am (and I'm course 6-1)

besides everyone should remember this from 8.02!

as far as this section goes, it's not just your coursework background. I know exactly what he's talking about and it was still a little hard to follow

Exactly. These are all just basic engineering concepts (OK, maybe not the UNIX). But it's still way more accessible than if he were using, say, drag equations from aero/astro or thermo from chemE.

I agree about the 6-3 part, but I think one of the problems with this is that "new" material from a completely different subject is being taught in just 2 pages.

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

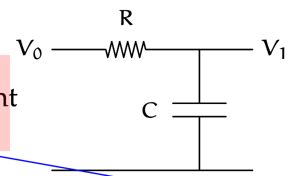

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \quad (2.11)$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \quad (2.12)$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

Don't capacitors and inductors make a circuit nonlinear?

This is what I thought too - he goes on to state that inductors and capacitors make the circuit have time dependent behavior and if I remember correctly an RC circuit with time is definitely not linear.

No, ideal resistors, capacitors, and inductors are part of linear circuits. Here linear means that when a linear combination of signals ax1(t) + bx2(t) is applied to the circuit F(x), the output is aF(x1)+bF(x2). Informally, what this means is that the *values of the electronic components*, (i.e. the resistance, cap, etc) dont change with voltage or current in the circuit.

Exactly, even if the voltage over tie, for example, is not linear, the *circuit* is still linear because its response to a sum of inputs is the same as the sum of its responses to those inputs.

I feel like it would be clarifying here to state what exactly it means when a circuit element has time depended behavior. It might be helpful to give a brief description of what each component does and how each varies with time. A resistor isn't time variant because the resistance is a fixed quantity, but it takes time for charge to be stored in a capacitor because the electrons have to be deposited.

I think its ok. He assumes basic knowledge

This sentence doesn't sound terribly grammatical, it might be the infinitive phrase being used as the subject. Perhaps a gerund prhase such as "Getting time-depending behavior..."

what's your definition of interesting here?

Interesting here means that the circuit's response changes with time. It's tough to do anything very useful with just resistors and voltage sources.

Exactly - resistors can provide us only with a linear response to voltage changes, whereas capacitors and inductors introduce some more interesting behavior as it takes time for them to respond to voltage changes.

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

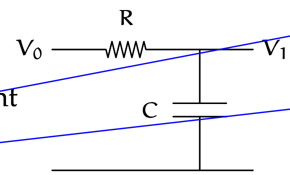

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

While you may assume that the reader understands what these elements are, it may be useful to either have a little note box clearly defining them - it may be nice for those who do know they are, but would like to be reminded. And from reading below and the comments, also including the governing equations for circuits and their elements will be useful when introducing (2.11)

side bars rock!

I agree that a reminder of their governing equations could help.

I think resistors are perfectly interesting, but perhaps this is like an electrical engineer saying levers are perfectly interesting.

and/or, since RLC circuits are good too.

Having an EE background I understand what this means, but a different audience might not catch that RC actually means resistors and capacitors.

Agreed. It would be very useful if you included a diagram of an RC circuit here.

I agree that it should say "an RC circuit, made up of a resistor and a capacitor," but isn't there a diagram right there?

write as V(t)

I think V0 here is constant, that's why he didn't write it as V(t)

I don't know the scope of your book, but maybe replace with "freshmen E&M class" to make it non-MIT friendly.

This equation is also covered a lot more in depth in 8.03 than in 8.02.

Unlike 8.03 though, 8.02 is a GIR.

Not that this is beyond my understanding, but everyone keeps expecting me to remember my GIRs. I wish someone had told me I'd need them after freshman year.

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

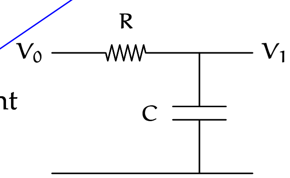

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \qquad (2.11)$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \qquad (2.12)$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

It's been a while since I've done this. It would help me a lot if you could justify this equation with a few arguments or convince me that it's correct. The dependence on the product RC is also not obvious. I guess this is what happens when your circuit E/M mostly was learned in high school...

I agree that equations might help us remember where this comes from (although that understanding isn't crucial to the rest of this analysis, it seems). The RC dependence is clear because the only time 'R' or 'C' appears is when the term 'R*C' appears. I found that to be well explained.

I'm not sure that the math needs to be explained - the chapter is about abstraction, and making RC = tau is very clear to me without all the mathematics that derived this differential equation. I think we're perfectly fine without it; it allows us to focus on the abstraction without getting bogged down in details.

i agree with this...skim the ee stuff, stick to the abstraction.

It would be really helpful if there was an explanation for these equations and this section in general. Since it has been a couple of years since I took 18.03 and 8.02, a lot of this material in this section is unfamiliar, and is therefore difficult to follow.

This seems to be a prime example of the conflict between not having enough explanation and being confusingly explicit. No matter what, someone is going to complain that it's either of the above. This example works fine using only the explanations given and vague memories of fulfilled GIRs. As mentioned before, see the tau substitution.

Maybe it would be useful to add an appendices for each of these sections. I think it's fine to use a birds eye view of things to ease clutter, but providing appendices would allow people reading the book to still follow in more detail.

kinda makes sense since they force each other. Although it seems a little strange cause shouldn't it just act like a capacitor if the resistor is really small

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

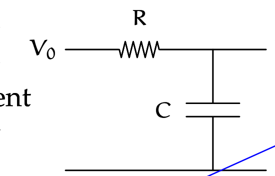

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

can you explain why you multiply the R and C together? what is the physical meaning of this product?

The sentence before explains it: it's a product that happens in the equation that has no importance as individual units, only as an abstraction.

It falls out of the math. The temperature analogy is more intuitive to explain. The time constant determines how quickly the temperature of a particular room, object etc, approaches that of a temperature it's exposed to. The higher this product the longer it takes. This makes sense - an ocean with a lot of capacity and resistance, takes a long time to change temperature.

This part is very clear to me. It is pretty easy to follow so far.

I agree this section was pretty easy to follow and useful, as it's been a while since I've dealt with circuits and needed a quick refresher.

i like this description–that it doesn't matter what R and C are, and that we don't need to know the detail to know what tau does

I agree, although this is a method I have often used for solving complex equations.

I agree, it seems like this is almost the opposite of an abstraction, that we are grouping together terms in order to ignore the confusion multiple variables could cause.

This explanation is a great example of abstraction, especially for those of us that have long forgotten the meaning of these equations.

What is this triple bar sign? Does it mean something like "from now on, let tau be defined as rc"?

Very clear abstraction! Some of the previous readings seemed to make them less clear.

Why is RC considered a time constant? Not familiar with circuits...

There wasn't much of a transition from defining t=RC to calling it a time constant. For the less EE savvy, is this confusing?

Yes, I think it should be brought up that it's a constant or a quantity with physical meaning, then that it is a TIME constant after the behavior is described.

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

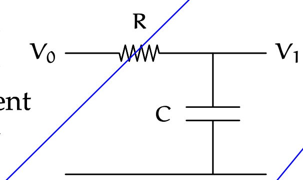

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V(1 - e^{-t/\tau}). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

My EE experience is limited, but is it possible for R or C to vary with voltage (ie the resistor becomes less resistive at high voltages)? Therefore, is it possible for tau to be dependent on V0 (or t, if V0 is a function)?

No, we are dealing with simple, fixed resistors and capacitors, which have a given resistance, R and capacitance, C respectively.

So... a step function? I think calling it by name might make this more familiar. Edit: it is actually called by name later... but I think it would be helpful to do so here (and perhaps still include the definition).

I felt this discription would be more familiar visually than saying step function, because it literally draws at the voltage, 0 and then to a constant in words as you read along.

Whoh, how did you get this again? Was this just a lucky guess from diff eq. intuition? Or is there some greater concept I'm missing?

i would appreciate an explanation also

All they did was take 2.11 and integrate it. They kept the dV1 term on the left and put dt and everything else on the right.

need to review 18.03. ew.

You just put dV to one side and dt to the other side of the equation, and then integrate. Very little 18.03 involved. (plus it is not that bad)

2.3 Low-pass filters

The next example is an analysis that originated in the study of circuits (Section 2.3.1). After those ontological bonds are snipped – once the subject is "considered independently of its original associations" – the core idea (the abstraction) will be useful in understanding diverse natural phenomena including temperature fluctuations (Section 2.3.2).

2.3.1 RC circuits

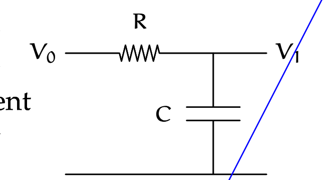

Linear circuits are composed of resistors, capacitors, and inductors. Resistors are the only time-independent circuit element. To get time-dependent behavior – in other words, to get any interesting behavior – requires inductors or capacitors. Here, as one of the simplest and most widely applicable circuits, we will analyze the behavior of an RC circuit.

The input signal is the voltage V_0, a function of time t. The input signal passes through the RC system and produces the output signal $V_1(t)$. The differential equation that describes the relation between V_0 and V_1 is (from 8.02)

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.11}$$

This equation contains R and C only as the product RC. Therefore, it doesn't matter what R and C individually are; only their product RC matters. Let's make an abstraction and define a quantity τ as $\tau \equiv RC$.

This time constant has a physical meaning. To see what it is, give the system the simplest nontrivial input: V_0, the input voltage, has been zero since forever; it suddenly becomes a constant V at $t = 0$; and it remains at that value forever ($t > 0$). What is the output voltage V_1? Until $t = 0$, the output is also zero. By inspection, you can check that the solution for $t \geqslant 0$ is

$$V_1 = V\left(1 - e^{-t/\tau}\right). \tag{2.12}$$

In other words, the output voltage exponentially approaches the input voltage. The rate of approach is determined by the time constant τ. In particular, after one time constant, the gap between the output and input

Another great explanation; this should be done for the majority of equations.

Agreed! This explanation makes the whole equation seem less intimidating and makes me feel like I don't need to search back into my 18.03/8.02/ old textbooks to understand what's going on.

I agree– the abstraction made with tau and the result here are very clear

I agree that the explanation of the equation is clear and very helpful, but I'm still unclear as to how we arrived at it. I guess if the point of the example is to show how abstraction is useful to understand the equation, then it doesn't really matter.

Definitely a good explanation. I wasn't really going to try to understand this equation on my own. Kinda glad they did that for me.

I agree this is a very helpful method, although I can't honestly say it's the first time I've ever seen it.

This is an easy example to explain, I think attempting to explain some previous equations would do more harm than help. Although I do agree I enjoyed seeing this step by step.

voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

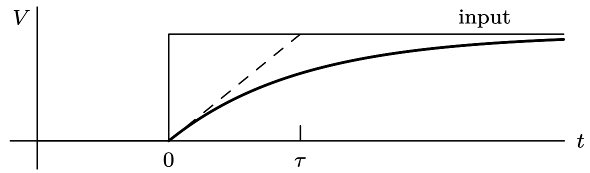

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.13}$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \tag{2.14}$$

or

$$B = \frac{A}{1 + i\omega\tau}. \tag{2.15}$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

is this the definition of a time constant?

i dont understand how this relates? why do we care about the dotted line?

I think the dotted line represents the output if the initial rate stayed constant, so I think it's there just to point out that voltage doesn't change at a constant rate and it actually decreases as afunction of time

Sorry I shouldn't have said "decreases as a function of time" but rather changes as one since it doesn't necessarily decrease.

Its just another interpretation of how tau relates to the transient response of the circuit

Yeah–I don't think it's relevant to solving the problem, but it might help some people understand the situation better and how tau fits in (part of an exponential rather than a linear model)

I agree that it gives real meaning to tau. "the time at which the voltage difference shrinks by a factor of e" is useful for determining tau but not necessarily intuitive, and doesn't offer much insight. I personally had never noticed the relation indicated by the dotted line, and it was definitely a cool thing to learn.

It was a new concept to me and a useful way of thinking about time constants. Is it only true for exponentials?

I'm a little confused by the graph itself. It would have been useful to have a line that represented what the output would be if the rate of approach remained its initial value.

voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

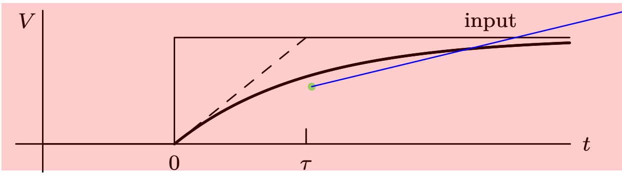

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.13}$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \tag{2.14}$$

or

$$B = \frac{A}{1 + i\omega\tau}. \tag{2.15}$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

You might consider marking the point on the V axis that the two lines converge upon.

Agreed, labeling it Vo would be useful for visualization

Vo=0... do you mean labeling it just 'V'? V1–>V as t–>+inf

Why is the time axis offset by so much from the 0 point? We aren't going to be considering negative time.

Otherwise the step rise from 0 to V falls right on the y axis and is then hard to distinguish from the y axis.

At least on my screen, it is still hard to discern a difference in the lines and if the voltage traces at t<0 are drawn on the X-axis.

I also found the fact that Tau is related to the initial rate, and the factor it decays by very interesting. Could you explain in the text the significance and reason for this relationship.

I think I should make the input and output graphs in black and the axix in gray (or not show it at all, since the t<0 region shows where the x axis is).

At this point, I'm wondering why this example is useful for this class. What was the point that was trying to be proven through this example?

I believe this example was more of a set up for the later example. In previous readings, examples have been given without any background knowledge, which confused many people. This reading actually takes its time to provide some sort of base when trying to understand the thermal problem. I found it very helpful.

Some example of really cool things with oscillatory inputs would be nice

Grammatically, this seems like an odd sentence

voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

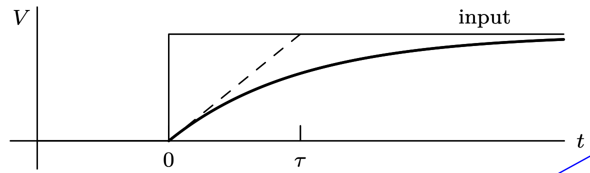

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.13}$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \tag{2.14}$$

or

$$B = \frac{A}{1 + i\omega\tau}. \tag{2.15}$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

I'd like an example or two?

I agree that another example would help, maybe one that is more EE focused.

I also agree. In all of my classes, we look at the step function because it's the easiest to learn from. It'd be nice to see a real world example of a sine wave, exponential, etc.

The example of daily temperature variations shows up in the next section, but there's no way to know that when you're reading this section. So it needs a mention right here.

Might be worth it to throw "Fourier" in here somewhere for people who might have seen it before.

I agree, it would have been helpful to see the Fourier here so we see the real part even if it is omitted.

But then people who don't know about Fourier stuff are going to complain that you put it in without explaining it. Dilemmas, eh? :P

But as written, I was somewhat confused about what it meant until I thought of the Fourier example

a. whats with the little side note things? b. don't complicate it unnecessarily

I think that the side note would be useful if it were actually put on the side...not in-line with the rest of the text

So (comma) let's...

This class really sneakily teaches you a lot of things you don't expect to learn. Not complaining, just observing.

voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

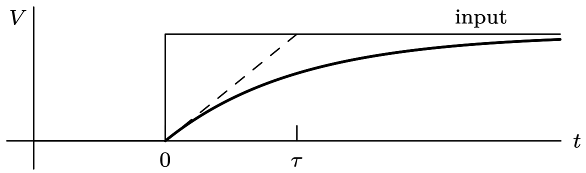

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.13}$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \tag{2.14}$$

or

$$B = \frac{A}{1 + i\omega\tau}. \tag{2.15}$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

What is "real part"? Is this a physics notation?

I think it's referring to the real part of a complex number (without the 'i' part)–since they said A is possibly a complex number

Here Ae^(iwt) = Asin(wt) + A*i*cos(wt), so the real part is A sin(wt).

I think it would be helpful to actually right exactly what real part you are referring to. Since for some this may be very new and abstract information, the more explicit you can make it the better

That's true when A is real. But if A is complex, e.g. x + yi, then the real part of Ae^(iwt) is x*sin(wt)-y*cos(wt).

Alternatively, if A is Ce^(i*phi), where C and phi (the phase) are real, then A*e^(iwt) is Ce^(i(wt+phi)), so the real part of A*e^(iwt) is C*sin(wt+phi).

After the confusion earlier since time-dependence was mentioned, this sentence really helped clear up in what way an RC circuit is linear. Maybe the fact that it is linear due to a linear diff eq in that paragraph near the beginning would help?

even with a oscillatory input?

You should probably have V_1(t) = Be^(iwt) written explicitly somewhere, and bpossibly show at least some of the algebra of substituting the 2 equations into 2.13

It could go right after "in the form", just as "V1 = Be..."

add tau in here to make it more coherent.

For me, the whether or not tau is in here doesn't make a difference since i know that tau is just RC. But it would help if there was a little more math shown–how we got to these expressions

Something (like RC=tau) needs to be different here, otherwise this is the same as eq. 2.11.

I agree that these are the same equation now and it would be nice to differentiate them

voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

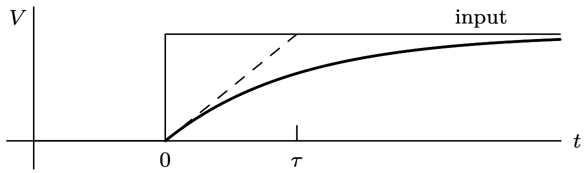

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.13}$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \tag{2.14}$$

or

$$B = \frac{A}{1 + i\omega\tau}. \tag{2.15}$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

I would like to see these math steps in a little more detailed manner. I am slightly confused what you are using for V1 and I would appreciate seeing the intermediate step before the common factor is removed.

I too was a little confused at first, but he actually defines what they are in the paragraph above. V1=Be^(iwt) and V0=Ae^(iwt).

Can we see the math for this simplification?

Also, this could be nice for the dimensional analysis leading to the quantity being unitless, as mentioned below by someone else.

Agreed - in this case, I think the math would be useful, and it doesn't take too much to put the values into the differential equation.

i'm not sure if you mean replacing V1 = Be^iwt and V0 = Ae^iwt and dV/dt = Biwte^iwt and then taking out the e^iwt but maybe just having a secondary step in the above equation which shows this replacement would be helpful. we can always go back and figure it out in our heads but it'd be nice just to see it anyway.

definitely, 18.03 was awhile ago for most of us.

Why not convert this to the conventional a + ib form, instead of having i in the denominator?

This form is much more applicable to the problem even if mathematically it isn't as pretty

why is this the transfer function? it does not seem to go along with what I have learned about them

Is it called a transfer function because it contains many generalizable points or is it that this particular function contains these points and there is another definition of a transfer function?

its a standard term. According to wikipedia: A transfer function (also known as the network function) is a mathematical representation, in terms of spatial or temporal frequency, of the relation between the input and output of a (linear time-invariant) system. " So to answer your question, yes

This is exactly what we're doing in 2.004 right now..

I like this way of breaking down the problem but how would this compare to the abstractions we've been doing, it seems like we're going in the opposite direction.

voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

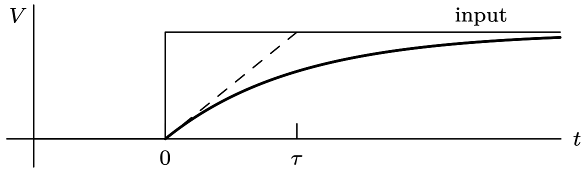

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \quad (2.13)$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \quad (2.14)$$

or

$$B = \frac{A}{1 + i\omega\tau}. \quad (2.15)$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

It might be useful to show the unit analysis here

Agreed.

Also agreed - not immediately clear that it's dimensionless.

angular frequency has units of 1/t and a time constant has units of t. Multiply and its dimensionless!

is there "credit" given for the checkbox "I agree" or do I really have to enter Agreed everywhere?

Neither gets credit! There's always a question from someone to answer, so you shouldn't ever feel like there's nothing to contribute.

Given a choice between clicking the "I agree" box or just typing in "I agree" as a comment, click the "I agree" checkbox. Then NB can tell me what comments are most pressing. It's not [yet] smart enough to parse the text, and uses only the explicit "I agree" button.

I agree that it should be clear that it's frequency*time and therefore unitless, but it's not that clear how R*C gives units of time.

Again, a great summary sentence!

Weird spacing. Seems like it should be one big paragraph since sentence transitions are "first, second, third."

I always find it useful when "large" is defined... If you go by the magnitudes of (1+iwt) vs. (iwt), then wt 10 gives you within 10%, I think.

Almost! It would be true without the "i". So, 1+wt is within 10% of wt, if wt 10. But with the "i", the mangitudes are within 0.5%.

Oops, I'm not sure what I did wrong. I tried to compare the magnitudes of (1+iwt) and (iwt).

Redoing it I get: |1+iwt|=sqrt(1+(wt)^2), and |iwt|=wt. At wt = 10, |1+iwt| 10.05, so there's your 0.5%.

Turns out (solving the algebra), you only need wt=2.2 (|1+iwt|=2.4) to get the magnitudes within 10%.

voltages shrinks by a factor of e. Alternatively, if the rate of approach remained its initial value, in one time constant the output would match the input (dotted line).

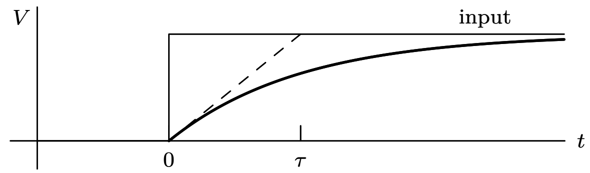

The actual inputs provided by the world are more complex than a step function. But many interesting real-world inputs are oscillatory (and it turns out that any input can be constructed by adding oscillatory inputs). So let's analyze the effect of an oscillatory input $V_0(t) = Ae^{i\omega t}$, where A is a (possibly complex) constant called the amplitude, and ω is the angular frequency of the oscillations. That complex-exponential notation really means that the voltage is the real part of $Ae^{i\omega t}$, but the 'real part' notation gets distracting if it is repeated in every equation, so traditionally it is omitted.

The RC system is linear – it is described by a linear differential equation – so the output will also oscillate with the same frequency ω. Therefore, write the output in the form $Be^{i\omega t}$, where B is a (possibly complex) constant. Then substitute V_0 and V_1 into the differential equation

$$\frac{dV_1}{dt} + \frac{V_1}{RC} = \frac{V_0}{RC}. \tag{2.13}$$

After removing a common factor of $e^{i\omega t}$, the result is

$$Bi\omega + \frac{B}{\tau} = \frac{A}{\tau}, \tag{2.14}$$

or

$$B = \frac{A}{1 + i\omega\tau}. \tag{2.15}$$

This equation – a so-called transfer function – contains many generalizable points. First, $\omega\tau$ is a dimensionless quantity. Second, when $\omega\tau$ is small and is therefore negligible compared to the 1 in the denominator, then $B \approx A$. In other words, the output almost exactly tracks the input.

Third, when $\omega\tau$ is large, then the 1 in the denominator is negligible, so

One great thing would be to summarize paragraphs or smaller sections with boxes similar to the way you summarize full chapters (or parts of chapters) with the little boxes. It makes it easier to get the exact point across; the first time through I missed a key point in the previous paragraph, and I think summary boxes (or bold type or something) may make the reader be attracted to the point,

$$B \approx \frac{A}{i\omega\tau}. \tag{2.16}$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by $90°$ relative to the input oscillations (where $360°$ is a full period). Why $90°$? In the complex plane, dividing by i is equivalent to rotating clockwise by $90°$. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

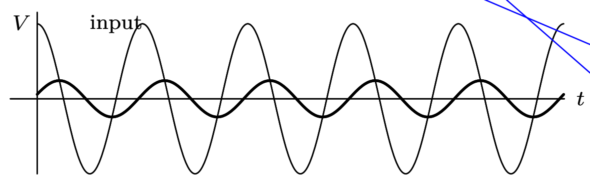

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

So I don't fully understand how this explanation (how the differential equation works) describes how a low pass filter works. I feel that the last paragraph accurately explains what a low pass filter does, but I dont understand from the previous explanation how it actually does this. It seems like one minute we are talking about the waveform and the differential equation and then we jump to a low pass filter.

Good Explanation.

I like this explanation. The page prior to this was important in understand how abstraction fits in but it's nice to get an explanation of some of the fuzzy points even though they may not be as relevant.

why does this happen? what leads to the 90 degree delay

you're going to explain this but not the equations or math?

Whenever we get to the symmetry part of the course, there are some nice examples in the 18.04 book that use symmetry to calculate long expressions of complex arithmetic (in particular powers of binomials)

Is this delay the interesting part that we figured out earlier, so we can apply it to more real life examples?

How exactly did you come to the conclusion that the delay is roughly 1 hour? Is it due to dividing by 90 (which is 1/4 of 360, so 1/4 of the input's 4 hour period) or is it due to wt=4 (so input's 4 hour period divided by wt is 1 hour)?

@ writer-of-comment: Your word choice is confusing. But the general idea is that the 1 hour is a quarter of the 4 hour period. Thus at that time, the max output should occur then.

I believe this is because 90/360 = 1/4 of a period. 1/4 of the 4-hour period gives us 1 hour after the input peaks for the output peaks.

Yeah, it's because it's a quarter of the 4 hour period–the graph right under it shows what he's talking about

yup 90/360 is 1/4th and 1 hour is 1/4th of 4. Would you ming showing us why it lags by 90 degrees.

$$B \approx \frac{A}{i\omega\tau}. \tag{2.16}$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by $90°$ relative to the input oscillations (where $360°$ is a full period). Why $90°$? In the complex plane, dividing by i is equivalent to rotating clockwise by $90°$. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

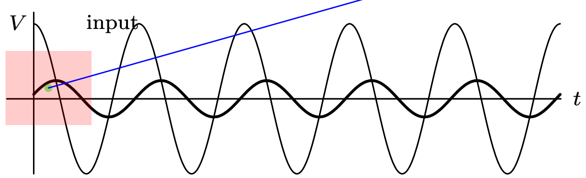

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

Is the output the bold curve in the graph? Why is it's amplitude smaller?

Its explained a little higher up, tat the amplitude by is shrunk by a factor of wt compared to the input amplitude

Again, it might help to label the V access

access? do you mean axis? in which case, it is labeled.

I would like to see a few labels on both axes to establish relative magnitudes.

If you just want relative magnitudes then can't those be deduced from the unlabeled graph? You can tell by looking that the taller fn is more than twice the amplitude, etc.

it'd be nice to label the two curves w/ a key or something, just for emphasis

I also agree, I think I get it but there's no need not to have labels.

I usually find your summary paragraphs helpful, but this one is a bit confusing. It kinda just goes over my head.

While I know that this is true, I'm not entirely this is an adequate follow up from the previous discussion. You prove that for omega*tau >>1 that the output is a dampened version of the input, but I can see a reasonable question being "so what if omega is really small (low frequency) but tau is just really big?" All I'm saying is perhaps you should add a line or two clarifying further that the "low frequencies" are not relative to normal frequency range but to the selected tau (that a large tau means only extremely low frequencies pass through and the rest are considered "high-frequency")

I'm not really sure that I follow how what we've done so far has demonstrated this.

I understand, but I think there's a missing bit that the limits of wt essentially mean limits of frequency–wt being small means low frequency, and wt large means high frequency. It could be confusing that there is still the factor of the time constant in there, but it just scales the exact values of "low" and "high" frequency. Later, in the thermal example, almost all frequencies are too high because it is the tau in that case that is usually high.

$$B \approx \frac{A}{i\omega\tau}. \quad (2.16)$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by $90°$ relative to the input oscillations (where $360°$ is a full period). Why $90°$? In the complex plane, dividing by i is equivalent to rotating clockwise by $90°$. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

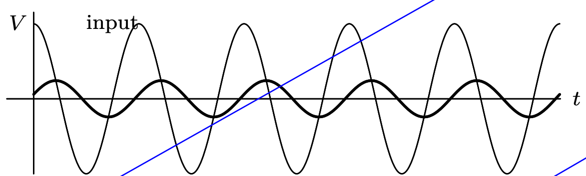

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

Great description of low- pass filters, simple and to the point.

Agreed, this may be the first time so far where enough background information was given so that lack of knowledge doesn't bog down the later example.

I agree... readers could actually skim, or not even understand much of what was talked about previously but understand its application here by this simple paragraph.

This is probably the best explanation of a low pass filter I have ever had

Was a circuit explanation of low pass filters really necessary? I took 6.003 and we learned about lowpass filters without using circuits- I was lost through the entire article until this paragraph

I don't understand the applications of this example.

It would be nice also if we can relate the low pass filter back to the beginning of the class where we talked about cd's and frequency. Low pass filters in electronic music to me is very interesting.

is the word "abstract[ed]" used here the same way as we've been talking about "abstraction" in general?

yeah, I think so. Because in the next section we go on to use this abstraction in another setting.

Yes it is, abstraction was defined in prevoius sections of Chapter 2, and this one is really just examples in different contexts.

It would be useful if throughout the text the information was put into an abstraction perspective.

Perhaps I'm just already used to seeing RC filters as examples of low pass filters, but I didn't quite follow how all this circuit analysis led to an abstraction away from the circuit origin.

I don't see how it is abstracted away from its origins

$$B \approx \frac{A}{i\omega\tau}. \tag{2.16}$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by $90°$ relative to the input oscillations (where $360°$ is a full period). Why $90°$? In the complex plane, dividing by i is equivalent to rotating clockwise by $90°$. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

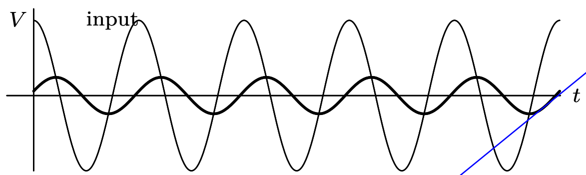

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

I thought it was really interesting that heat flow is so closely connected with RC circuits. But I would have liked a short explanation of why these two concepts are so closely connected. I feel like, in this section, I learned how to reappropriate the RC circuit equations for use as heat flow equations, but I didn't know why I was allowed to do so.

I agree.

I guess looking back, not knowing what resistors and capacitors do physically/qualitatively makes the thermal analog not as clear. A brief explanation as to how capacitors store charge and how resistors dampen current flow might make the transisition from circuits to thermodynamics more clear.

This idea is pretty interesting to me. Having seen a lot of stuff related to RC circuits through course 6 classes I'm always amazed at how many things they can be used to model.

This might make the section too lengthy; it's to serve a purpose for the section; if someone is then further interested on capacitors, they should look it up. But that's just my point of view; I like these chapters short.

It would be nice if we could go over the "pool temperature" exercise in class. I somehow don't quite get how to solve it correctly.

$$B \approx \frac{A}{i\omega\tau}. \quad (2.16)$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by $90°$ relative to the input oscillations (where $360°$ is a full period). Why $90°$? In the complex plane, dividing by i is equivalent to rotating clockwise by $90°$. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

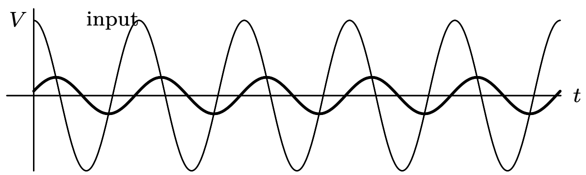

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

Where can we see an example of a high-pass filter? And how do we know that this is a low pass filter?

Heat flow can be understood as a low pass filter because the time constants of heat flow systems are usually high. That is, it takes a relatively long time for the system to respond to changes in temperature. And if the temperature is changing very rapidly (high frequency), then the system will respond minimally (the high frequency fluctuations will be filtered out).

Yup. It makes more sense if you've actually seen a high pass and low pass filter work, because you can figure out what they do to the frequency, and then apply the concepts in real life. I think it would help a little if there was a demo involved with this section, but that is hard to do on paper!

"Rapidly" and "slowly" are relative terms. Looking at the value of w*t will give a clue to whether we are in the high or low frequency regime. If w*t is large, then the frequency is relatively high for this time constant. If w*t is low, then the frequency is relatively low.

Heat transfer behaves like a low pass filter because it can be modeled by the same RC circuit as before (heat capacity is like capacitance, and thermal resistance is like electrical resistance). Any system that can be modeled in such a way will act as a low pass filter, regardless of the time constant. The value of the time constant only affects what counts as a "high" or "low" frequency.

Thanks for the clarification I would've had the same question.

that makes complete sense and i've never thought of it before, not in 2.003 or 2.005 or 8.02

how is it a model for heat flow?

this makes sense because the heat flow through materials has resistance as well as the material holding energy

$$B \approx \frac{A}{i\omega\tau}. \tag{2.16}$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by $90°$ relative to the input oscillations (where $360°$ is a full period). Why $90°$? In the complex plane, dividing by i is equivalent to rotating clockwise by $90°$. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

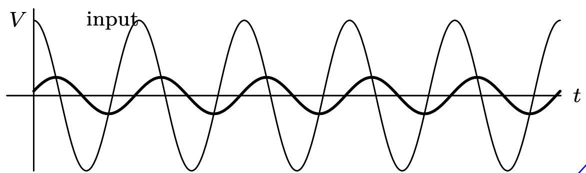

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

Interesting point! I never really thought of the connection before, but it makes sense after reading this paragraph.

I agree. I really liked this paragraph and the interesting insight it brought up.

This is a really cool comparison. I never thought about this before, and the way you connected the more difficult low-pass filter with the everyday cup of coffee makes both concepts easier to understand.

If anyone is interested in this, 18.303 goes into it a little more in depth.

i dont think you're referring to this example, which is extremely simple, but everyone course 2 here has taken thermo.

I guess this is why in Coffee shops they warm the cups and jars on the surface above the machine, so that they delay the time constant and the milk and coffee is kept cold for longer periods of time.

Perhaps this example could be placed before the RC circuits as it is easier to relate to.

I think it's best how it is, since this one is easier to just relate to the RC circuit. Also, logically it makes sense because here we see how the abstraction of resistance and capacitance actually apply to other seemingly unrelated situations–like a cup of tea cooling down–which is pretty crazy

this was a very good explanation relating the electrical circuit model to the thermal model

perhaps i'm being too nit-picky here, but i believe that, technically, heat cannot be "stored" because it is defined to be heat transfer.

thermal energy...?

This is a good description of the components of the system in terms of their electrical analogies.

do you really need to say "so-called"? we get it.

Perhaps he is defining a difference between the term thermal capacitance as a way of comparing the behavior with an EE capacitor and the actual physical behavior of the process in question

$$B \approx \frac{A}{i\omega\tau}. \tag{2.16}$$

In this limit, the output variation (the amplitude B) is shrunk by a factor of $\omega\tau$ in comparison to the input variation (the amplitude A). Furthermore, because of the i in the denominator, the output oscillations are delayed by $90°$ relative to the input oscillations (where $360°$ is a full period). Why $90°$? In the complex plane, dividing by i is equivalent to rotating clockwise by $90°$. As an example of this delay, if $\omega\tau \gg 1$ and the input voltage oscillates with a period of 4 hr, then the output voltage peaks roughly 1 hr after the input peaks. Here is an example with $\omega\tau = 4$:

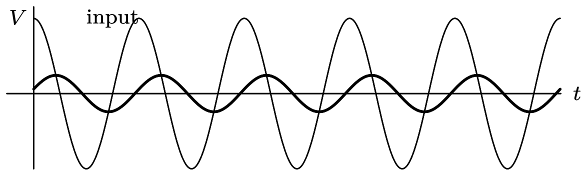

In summary, this circuit allows low-frequency inputs to pass through to the output almost unchanged, and it attenuates high-frequency inputs. It is called a low-pass filter: It passes low frequencies and blocks high frequencies. The idea of a low-pass filter, now that we have abstracted it away from its origin in circuit analysis, has many applications.

2.3.2 Temperature fluctuations

The abstraction of a low-pass filter resulting from the solutions to the RC differential equation are transferable. The RC circuit is, it turns out, a model for heat flow; therefore, heat flow, which is everywhere, can be understood by using low-pass filters. As an example, I often prepare a cup of tea but forget to drink it while it is hot. Slowly it cools toward room temperature and therefore becomes undrinkable. If I neglect the cup for still longer – often it spends the night in the microwave, where I forgot it – it warms and cools with the room (for example, it will cool at night as the house cools). A simple model of its heating and cooling is that heat flows in and out through the walls of the mug: the so-called thermal resistance. The heat is stored in the water and mug, which form a heat reservoir: the so-called thermal capacitance. Resistance and capacitance are transferable abstractions.

but this isn't what the abstraction was in the previous section. in the low pass filters we used tau as the abstraction instead.

I agree with that except for the "instead". Tau was one of the abstractions, and resistance and capacitance are two further abstractions.

This makes sense to me.

I never thought of the analysis for heat before, very interesting

thats really cool! haven't taken any thermo class since I'm 6-1, but the analagy with RC really resonates.

I like the last sentence in this paragraph - it summarizes everything nicely and drives home the purpose of the previous section.

Can you show this mathematically?

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

Comments on page 4

This is a very neat way of looking at thermal transfers. However, I am not sure how the diagram for a low-pass filter translates to this cup of tea. It's just non-obvious how a the thermal capacitor or thermal resistor should be connected with thermal wires and where the temperatures should be sampled.

I'm not really sure but I think that these two things are juxtaposed due to the time constant similarity

Oh, that's a really clever analog! I took thermo last term, but we never approached it in this manner.

Really? In 2.005 we did, and I just realized now that my thought process in solving those problems was abstracting things – treating physical objects as resisters and capacitors to analyze them. Neat that this chapter helped me explicitly realize that now.

its amazing how many things you can model with resistance and capacitance. I think we did something in 2.004 where we modeled springs and dampers similarly... yay cross-disciplinary connections :)

Yes! One thing I've really liked about this class is, for these kinds of problems, it seems to either explain what i've been doing in my head or give me some new wayt o approach problem!

Could this alternatively be calculated using Newton's Law of Cooling?

This is a very clear paragraph, the time constant example is explained well here

I agree. This paragraph strikes a nice balance of technical and understandable.

I particularly liked the clear explanation of how one might go about estimating tau for a real system.

temperature estimation of lukewarm...?

oh, it's answered below

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

Isn't this not true, since F is not on an absolute scale from 0?

I agree. I don't really understand how you could make this calculation accurate. If you change units (to Kelvin, let's say) you would definitely not get a factor of 4. Farenheit is a pretty arbitrary unit....

That flaw in the Fahrenheit scale doesn't matter here since it's just the delta T that matters. The same gap between room temp. and tea temp. might have gone from 55C-21C = 34C when hot to 30C - 21C = 9C when it cooled, or also about a factor of 4 (with the differences just scaled by the 9/5 F/C ratio)

Is there a quality factor Q for this, as there is for an RLC circuit?

LRC circuits are second order circuits whereas RC circuits are 1st order circuits. Only seconds order circuits have Q factors.

Wouldn't this also depend on surface area and other variables that aren't used to calculate time constant?

Is this saying we can find tau by letting exp(-t/tau) = 2.72?

Close, I think it's that exp(-t/tau) = 1/2.72. If you solve that equation by taking ln of both sides, you get t=tau.

It would have been helpful to see this quick comparison when tau was initially introduced.

why does the concept of time constant come from an RC circuit? I think if you switched the order of these examples it would make more sense

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

Why is the frequency day^-1? Are we just considering the period to be 1 day? And ignoring day-to-day changes? (a summer day being much different than a winter day)

I think that's the frequency of the driving function, the daily fluctuation of temperature. However, the RC circuit does not change the overall frequency (see previous page), so the overall frequency is the frequency of the driving function.

I believe the assumption is that each day's temperature fluctuation is, more or less, the same. So the room's temperature at a given time of day is the same as it is +-24 hours.

I agree with the above comment - I think the idea is the temperature at 9am today will be the same as the temperature at 9am tomorrow, giving you the 1-day period, or the 1-day^-1 frequency.

but frequency in this context is a little weird.

Just think of it as 1/period as mentioned above

This is also confusing because it doesn't specify that we're just leaving a teacup in a room for a time of a day or so AFTER it has come to equilibrium.

Are the sub-brackets necessary?

I think so.

Agreed, I think they are necessary. I probably would have been confused without them, or I would have had to continually scroll back up to where the equation was initially written. I don't think they detract at all from the text or equation.

I would definitely say yes, they make the point clear for us and the analog to RC circuits easy to understand.

If they are necessary, the only formatting thing that seems to bug me (and it may just involve being OCD about alignment) is that the f sub-bracket dropped below the others

I am finicky about alignment as well. I wrote a few TeX functions (macros) to handle that; I'll insert the magic code and make the alignment work out.

This whole section seemed clear and easy to follow. It might be kind of nice to relate it to Fourier's law of conduction.

Could this be explained quantitatively by any chance?

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

so what are the possible factors that affect the thermal resistance? length, surface area, and what else?

here we're looking at thickness of the walls and surface area–the thicker the wall, more thermal resistance. but at the same time, teh greater the surface area, the more heat is able to escape, so smaller thermal resistance

area, material

Are we just ignoring the fact that a teacup has no top? Would it be fair to assess the comparable thermal resistance between a teacup and a house with no roof?

Good point for a further refinement, though we aren't directly comparing the two physical situations, just their overall time constants (abstraction). It's more that in the first example we've neglected convection off the top of the liquid, which *is* a significant route of heat transfer from a teacup.

edit to my earlier note: We didn't really neglect convection but we lumped it into our abstraction, since we didn't consider modes of heat transfer, just our empirical experience of the time constant as our tea cooled.

http://en.wikipedia.org/wiki/R-value_%28insulation%29 Thermal conductivity is covered here.

although the house has more surface area and will lose more heat to convection - conduction is a much more efficient heat dispenser and the resistance is greater in the house (due to it's material properties)

never thought of that

I don't entirely understand what is compensating for what here. But besides this, I really like this example, it really makes the point of the article clear.

So if the house has a smaller thermal resistance, shouldn't it have a SMALLER time constant?

So I didn't do UPOP, but some of my friend's did, so this question goes out to you guys. Does a house really have smaller thermal resistance than a cup of tea?

this is how they explain heat transfer in 005- with resistances

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

This paragraph helped me a lot in understanding how the house compares both to the tea cup and a circuit.

I think I may be getting the results of resistance and capacitance confused here in terms of the time constant. There is a lower thermal resistance (smaller R) but it is balanced by a significantly larger C, so the overall time constant is greater, right?

so, just to clarify, thermal capacitance is related to the size/volume of the object?

not quite; heat capacity has the units of J/kg-K, so if there is more mass then there is more thermal capacitance. you can have a larger volume but less thermal capacitance if the materials are different.

specific heat capacity is J/kg-K, which you multiply by mass to get heat capacity of, e.g., your brick house.

So the air mass is the capacitance and the walls the resistors to temperature difference, right? So this can be applied to anything that stores energy and looses it over time.

The walls are both resistance and a contributor to the capacitance. I suspect most of the thermal capacitance (often called thermal mass) is in the walls and other solid objects, rather than in the air.

When we study random walks, we'll be able to estimate the relative contributions. The paper I was reading about time constants of Greek houses said that furniture made a significant contribution to the time constant. I think it was 30 hours (out of the total of 90 hours).

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

I'm a little confused on how a house can have a 4 day time constant when each day the outside temperature rises and falls...thus making temperature fluctuations inside of the house a daily thing

Ah, don't confuse the time constant, which is an intrinsic property of the house, with the input or driving frequency omega (or with 1/omega). They can be totally different.

Because the system is linear, the driving frequency is also the output frequency, and neither need be the same as 1/tau. Rather, the time constant helps you figure out the output amplitude given the input amplitude (e.g. if the input frequency is much higher than 1/tau, then the output amplitude is small compared to the input amplitude).

If it's anything like mine, your Cambridge/Somerville apartment probably has a time constant closer to South African than Greek. You can tell because the delay is short between when it gets cold or hot outside and your apartment follows suit.

this was also confusing me but the response helped. However, I wonder how you test the time constant of a house if you don't have conditions where you can heat it up and then let it drop back down (like the tea cup) without it being greatly affected by outside conditions

why Greece?

Greece gets quite hot in the summers and thus it is necessary to have well insulated houses to avoid getting baked in the summers. Also, older house construction methods in the Mediterranian region were designed to be well insulated (narrow streets, houses built with thick clay/stone walls, etc)

Does the temperature of the room affect the time constant?

In a linear system, the time constant is independent of the temperature (just as the RC time constant is independent of the voltages). No system is completely linear, but both the room and the RC circuit are very good approximations.

If R_t is the thermal resistance and C_t is the thermal capacitance, their product R_tC_t is, by analogy with the RC circuit, a thermal time constant τ. To measure it, heat up a mug of tea and watch how the temperature falls toward room temperature. The time for the temperature gap to fall by a factor of e is the time constant τ. In my extensive experience of neglecting cups of tea, in 0.5 hr an enjoyably hot cup of tea becomes lukewarm. To give concrete temperatures to it, 'enjoyably warm' is perhaps 130 °F, room temperature is 70 °F, and lukewarm is perhaps 85 °F. The temperature gap between the tea and the room started at 60 °F and fell to 15 °F – a factor of 4 decrease. It might have required 0.3 hr to have fallen by a factor of e (roughly 2.72). This time is the time constant.

How does the teacup respond to daily temperature variations? In this system, the input signal is the room's temperature; it varies with a frequency of $f = 1\,\text{day}^{-1}$. The output signal is the tea's temperature. The dimensionless parameter $\omega\tau$ is, using $\omega = 2\pi f$, given by

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.3\,\text{hr}}_{\tau} \times \frac{1\,\text{day}}{24\,\text{hr}}, \tag{2.17}$$

or approximately 0.1. In other words, the system is driven slowly (ω is not large enough to make $\omega\tau$ near 1), so slowly that the inside temperature almost exactly follows the outside temperature.

A situation showing the opposite extreme of behavior is the response of a house to daily temperature variations. House walls are thicker than teacup walls. Because thermal resistance, like electrical resistance, is proportional to length, the house walls give the house a large thermal resistance. However, the larger surface area of the house compared to the teacup more than compensates for the wall thickness, giving the house a smaller overall thermal resistance. Compared to the teacup, the house has a much, much higher mass and much higher thermal capacitance. The resulting time constant R_tC_t is much longer for the house than for the teacup. One study of houses in Greece quotes 86 hr or roughly 4 days as the thermal time constant. That time constant must be for a well insulated house.

In Cape Town, South Africa, where the weather is mostly warm and houses are often not heated even in the winter, the badly insulated house in which I lived had a thermal time constant of around 0.5 day. The dimensionless parameter $\omega\tau$ is then

Just wondering, does humidity play a role? Warm air with more moisture having a longer or shorter time constant?

That is a very interesting question. I don't know the answer but I suspect that humidity leads to greater thermal transfer and thus reduces the thermal time constant. Heat is transfered through phonons (vibrations of molecules) and higher humidity means higher water vapor and thus more efficient heat transfer.

But humid air has a higher specific heat, though the effect is small.

It might be nice to give the dimensionless time quantity for the houses in Greece. Not that it's hard to find, but it'd just be more convenient to have direct comparisons between the teacup, the house in Greece, and the house in South Africa.

This sentence is hanging out, and read funny. I think something more similar to, "That must be a well insulated house!"

I agree - I was confused as to how you were able to assert that the time constant must be for a well insulated house. Is there a reason homes in Greece are "well insulated"? I like the inclusion of the South African homes later, but this sentence really just sticks out awkwardly

These facts are pretty interesting.

Why were you living in South Africa?

I was one of the founding faculty of the African Institute for Mathematical Sciences (http://www.aims.ac.za) in Cape Town, and I taught several of the courses.

And, almost all my better half's family is still in South Africa.

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.5\,\text{day}}_{\tau}, \tag{2.18}$$

or approximately 3. In the (South African) winter, the outside temperature varied between 45 °F and 75 °F. This 30 °F outside variation gets shrunk by a factor of 3, giving an inside variation of 10 °F. This variation occurred around the average outside temperature of 60 °F, so the inside temperature varied between 55 °F and 65 °F. Furthermore, if the coldest outside temperature is at midnight, the coldest inside temperature is delayed by almost 6 hr (the one-quarter-period delay). Indeed, the house did feel coldest early in the morning, just as I was getting up – as predicted by this simple model of heat flow that is based on a circuit-analysis abstraction.

I think that while the Unix example in the previous section was more interesting to the course 6 audience, this example is more interesting and understandable for the course 2 audience.

I really like this example. I did not think I would ever be able to understand the pool changing temperature in the diagnostic and problem set, but now i do.

Same here. The house temperature variations is a really good example of using extraction to solve a problem.

I would've thought the coldest temperature is near dawn, since the outside air had all night to cool.

this makes sense to me i just get a little confused when it cycles to a new day

So this really is just a universal truth?

I'm sure there are varations and this is a massive oversimplification, but it does make sense by experience, no?

Agreed - I'm sure it can't be true for everything, but that does sound about right for my house, and it's probably just a general rule of thumb to follow.

A little intuition tells us that 1/4 period is reasonable. More than that would mean that your house would still be cooling off later in the morning even after the temperature outside had begun to warm back up for a while.

Previously this was said to apply when omega*tau>>1...but I wouldn't consider 3>>1, so perhaps this should be clarified or the restriction relaxed to just >1?

slightly confused. 0.5 is the period which is 12 hours, right? so isn't 1/4 of the period 3 hours, not six? (nevermind i got it, the time constant is 0.5 but the period is 1]

Finally someone explains to me why my dorm room is so cold when I wake up even though it is warm when I go to bed at the coldest point in the night.

This is also a little biased by presumably being covered in blankets for most of the night...

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.5\,\text{day}}_{\tau}, \tag{2.18}$$

or approximately 3. In the (South African) winter, the outside temperature varied between 45 °F and 75 °F. This 30 °F outside variation gets shrunk by a factor of 3, giving an inside variation of 10 °F. This variation occurred around the average outside temperature of 60 °F, so the inside temperature varied between 55 °F and 65 °F. Furthermore, if the coldest outside temperature is at midnight, the coldest inside temperature is delayed by almost 6 hr (the one-quarter-period delay). Indeed, the house did feel coldest early in the morning, just as I was getting up – as predicted by this simple model of heat flow that is based on a circuit-analysis abstraction.

Your blowing my mind Professor Sanjoy.

Agreed, its awesome seeing how these simple abstractions and comparisons let you analyze real world situations quickly.

Is it just a coincidence that the RC abstraction could apply to heat flow, or could we find similar abstractions throughout the physical sciences?

The RC abstraction can be applied to _many_ different systems...think back to any physics class you've taken that used /exp/ ever...same concept.

I would imagine it would apply to anything which can be modeled as a value flowing through a system. So it probably extends to fluid dynamics, and I'd guess air flow as well. Of course, all of these would only be rough approximations. Fortunately, that's what we're looking for in most cases.

I like the example used in this section. While following the math and circuitry may have been a little tedious at first, its application to simple heat flow within the home was very unique and easy to follow!

I agree. I have seen circuits as models for various systems and it really emphasizes abstraction.

Just checking, but abstraction is reusing principles of something we know on something new. But all this approximation seems like a stretch. Is this not so?

I really liked the variety of examples in this section. Something more typically analytical and "science-y," an everyday application, and a broadened application that helps explain a larger aspect of life.

I very much liked the examples included here. The EE example was perhaps a little too in depth, but the thermal examples made it much clearer.

I think that the variety of examples not only helps to ensure that more people can relate to at least one example and gain a better understanding, but it provides some room to compare and contrast so that we can pick out the similarities between the examples and see what information is applicable to different situations and what was specific to one example.

I happened to also like the thermal example as opposed to the circuit one. ME vs EE for you.

$$\underbrace{2\pi f}_{\omega}\,\tau = 2\pi \times \underbrace{1\,\text{day}^{-1}}_{f} \times \underbrace{0.5\,\text{day}}_{\tau}, \tag{2.18}$$

or approximately 3. In the (South African) winter, the outside temperature varied between 45 °F and 75 °F. This 30 °F outside variation gets shrunk by a factor of 3, giving an inside variation of 10 °F. This variation occurred around the average outside temperature of 60 °F, so the inside temperature varied between 55 °F and 65 °F. Furthermore, if the coldest outside temperature is at midnight, the coldest inside temperature is delayed by almost 6 hr (the one-quarter-period delay). Indeed, the house did feel coldest early in the morning, just as I was getting up – as predicted by this simple model of heat flow that is based on a circuit-analysis abstraction.

Good section, with the right amounts of explanation and interesting information.

3
Symmetry and conservation

3.1	Heat flow	46
3.2	Cube solitaire	48
3.3	Drag using conservation of energy	50
3.4	Cycling	54
3.5	Flight	56

Symmetry greatly simplifies any problem to which it applies – without any cost in accuracy. A classic example is the following story about the young Carl Friedrich Gauss. The story might be merely a legend, but it is so instructive that it ought to be true. One day when Gauss was 3 years old, the story goes, his schoolteacher wanted to occupy the students for a good while. He therefore asked them to compute the sum

$$S = 1 + 2 + 3 + \cdots + 100,$$

and then sat back to enjoy a welcome break. To the teacher's surprise, Gauss returned in a few minutes claiming that the sum is 5050. Was he right? If so, how did he compute the sum so quickly?

Gauss noticed that the sum remains unchanged when the terms are added backward from highest to lowest. In other words,

$$S' = 100 + 99 + 98 + \cdots + 1$$

equals S. Then Gauss added the two sums:

$$\begin{array}{r} S = 1 + 2 + 3 + \cdots + 100 \\ +S = 100 + 99 + 98 + \cdots + 1 \\ \hline 2S = 101 + 101 + \cdots + 101. \end{array}$$

i thought this was the most interesting (and my favorite) reading so far.

i still really hate this webpage. why do we have to click on each new page, and why does it always have to adjust to the top of the page when you do that? it definitely inhibits concentration on the actual material, which you think would be the important part of the class. also, if you click anywhere a stupid little box pops up and blocks everything you're trying to read or type.

I agree that having to remember to click on each new page as you come to it is irritating. It would be nice if NB would realize that once you'd scrolled a bit down a page you were most likely reading that page, instead of the last page you clicked on.

You can always look at the pdf in your own reader and go back to nb to look at and enter comments.

I agree about the having to click on each new page. I have been trying to convince Sacha that this is a problem. He has probably just been too busy to hear me on this point.

But, I'll forward this thread to him and see if the feedback from other users is enough to convince him.

When do you know to use symmetry as a abstraction versus other methods we've learned? Is symmetry good only for simple summation problems, etc?

Symmetry is such an amazing way to solve problems. I'm course 8, and a lot of professors make a really big deal about symmetry arguments. This kind of logic has become intuitive to me and there are so many times where I can solve a problem just by looking at it now.

This works for the centers, but what about some random point in the pentagon area? and what about prior to equlibrium. Gaussian symmetry addition has very limited uses.

I really like the thermo example. It applies the message to something I understand (being course 2).

I'm confused, are the temperatures 120 degrees or are they 24 degrees at the center?

I didn't find any problems with this reading; the examples were very straightforward and easy to follow. Even non-technical people can follow along and grasp the concept.

3

Symmetry and conservation

3.1 Heat flow 46
3.2 Cube solitaire 48
3.3 Drag using conservation of energy 50
3.4 Cycling 54
3.5 Flight 56

Symmetry greatly simplifies any problem to which it applies – without any cost in accuracy. A classic example is the following story about the young Carl Friedrich Gauss. The story might be merely a legend, but it is so instructive that it ought to be true. One day when Gauss was 3 years old, the story goes, his schoolteacher wanted to occupy the students for a good while. He therefore asked them to compute the sum

$$S = 1 + 2 + 3 + \cdots + 100,$$

and then sat back to enjoy a welcome break. To the teacher's surprise, Gauss returned in a few minutes claiming that the sum is 5050. Was he right? If so, how did he compute the sum so quickly?

Gauss noticed that the sum remains unchanged when the terms are added backward from highest to lowest. In other words,

$$S' = 100 + 99 + 98 + \cdots + 1$$

equals S. Then Gauss added the two sums:

$$\begin{array}{r} S = 1 + 2 + 3 + \cdots + 100 \\ +S = 100 + 99 + 98 + \cdots + 1 \\ \hline 2S = 101 + 101 + \cdots + 101. \end{array}$$

Read the introduction and first section (3.1) for a reading memo due Thurs at 10pm. This is the first part of a new unit on symmetry and conservation, itself part of a broader unit on "Lossless methods of discarding complexity".

Is symmetry the same thing as abstraction? Because the example below looks exactly like what we did in the last pset.

Symmetry can be considered a form of abstraction; since abstraction involves a resuable, 'modular' approach, this is in fact a similar technique.

Is this how you comment?

This is an excellent story to begin this chapter on symmetry. I think it does a good job of setting the tone for the rest of the sections in this chapter.

what is he ended up doing as a career?

He went on to become one of the world's greatest mathmaticians.

Gaussian surfaces!

Gaussian distribution - that's him!

3? I wouldn't have thought of doing this now.

I've heard this story before, but I'm pretty sure the version I've seen just said he was in school, I don't think it made any claims about his age. But it doesn't seem likely that children would have started school before age 5 or 6.

I agree. Back in the day, kids also started school later on.

3 years old definitely seems a bit of a stretch...I was probably learning how to count (if even), though considering he was a genius mathematician, it could have been 3.

Agree. Maybe the story could be fact-checked, or if nothing else the age bumped up to something like 6 so it's a bit more believable.

Yea, 3 years old is a bit disheartening.

3
Symmetry and conservation

3.1	Heat flow	46
3.2	Cube solitaire	48
3.3	Drag using conservation of energy	50
3.4	Cycling	54
3.5	Flight	56

Symmetry greatly simplifies any problem to which it applies – without any cost in accuracy. A classic example is the following story about the young Carl Friedrich Gauss. The story might be merely a legend, but it is so instructive that it ought to be true. One day when Gauss was 3 years old, the story goes, his schoolteacher wanted to occupy the students for a good while. He therefore asked them to compute the sum

$$S = 1 + 2 + 3 + \cdots + 100,$$

and then sat back to enjoy a welcome break. To the teacher's surprise, Gauss returned in a few minutes claiming that the sum is 5050. Was he right? If so, how did he compute the sum so quickly?

Gauss noticed that the sum remains unchanged when the terms are added backward from highest to lowest. In other words,

$$S' = 100 + 99 + 98 + \cdots + 1$$

equals S. Then Gauss added the two sums:

$$\begin{array}{r} S = 1 + 2 + 3 + \cdots + 100 \\ +S = 100 + 99 + 98 + \cdots + 1 \\ \hline 2S = 101 + 101 + \cdots + 101. \end{array}$$

I disagree with this phrase. There are plenty of things out in the world that are "instructive" but cannot be true.

Yeah I found this phrasing very strange too...it's quite awkward.

I think it's rather amusing, he's trying to say that something should be true because it proves a very good point, and thus as a true story would be exemplary.

I like it. I'm with (9:44)

I'm with 9:44 also, fact and fiction are often constructed out of what we want to believe. Especially with history, so this statement was pretty amusing to me too.

3 years old!?!?!

I know–so crazy. I'm reading about this method now and i have to think about it for a sec before I'm like wow thats really clever

i drew a plan view drawing when i was 3 to show my architect mother, and she thought that was cool. i've always been better with art than numbers though.

That is very impressive...

Good example, remember that from junior high!

It seems like all the greatest scientific minds have stories like this. I wonder how many people at MIT have a story like this about themselves?

I don't think this phrase is actually needed here.

I used this in the last homework

I wish I saw this for the homework it would have been useful

I've seen this during an interview before..pretty cool how Gauss figured it out when he was so young!

The story is likely apocryphal, but the idea is clever nonetheless. Remarkable how many of these questions show up during interviews!

3
Symmetry and conservation

3.1	Heat flow	46
3.2	Cube solitaire	48
3.3	Drag using conservation of energy	50
3.4	Cycling	54
3.5	Flight	56

Symmetry greatly simplifies any problem to which it applies – without any cost in accuracy. A classic example is the following story about the young Carl Friedrich Gauss. The story might be merely a legend, but it is so instructive that it ought to be true. One day when Gauss was 3 years old, the story goes, his schoolteacher wanted to occupy the students for a good while. He therefore asked them to compute the sum

$$S = 1 + 2 + 3 + \cdots + 100,$$

and then sat back to enjoy a welcome break. To the teacher's surprise, Gauss returned in a few minutes claiming that the sum is 5050. Was he right? If so, how did he compute the sum so quickly?

Gauss noticed that the sum remains unchanged when the terms are added backward from highest to lowest. In other words,

$$S' = 100 + 99 + 98 + \cdots + 1$$

equals S. Then Gauss added the two sums:

$$\begin{aligned} S &= 1 + 2 + 3 + \cdots + 100 \\ +S &= 100 + 99 + 98 + \cdots + 1 \\ \hline 2S &= 101 + 101 + \cdots + 101. \end{aligned}$$

Ah yes I remember this story from elementary school.

It's definitely amazing how taking a step back and looking at a problem in a new way can help so much. I often find myself trying to forget things I know when I look at this type of problem to catch some form of symmetry that would help me out

I really love this story and I actually use the idea all the time to solve problems. In fact, I rarely bother to remember what the formula is for the sum of sequential numbers, sum of sequential even or odd numbers, etc. because I know that I can always derive it again from this trick.

Exactly. And then sit and feel proud about myself for deriving it until I remember how young Gauss was when he figured it out, and the fact that I can only derive it because I know the trick.

edit: "sum *was* 5050"

I've seen this example before but never thought of it as a "symmetry" problem. It's obvious now, and this example will help me find future solutions.

3

Symmetry and conservation

3.1 Heat flow 46
3.2 Cube solitaire 48
3.3 Drag using conservation of energy 50
3.4 Cycling 54
3.5 Flight 56

Symmetry greatly simplifies any problem to which it applies – without any cost in accuracy. A classic example is the following story about the young Carl Friedrich Gauss. The story might be merely a legend, but it is so instructive that it ought to be true. One day when Gauss was 3 years old, the story goes, his schoolteacher wanted to occupy the students for a good while. He therefore asked them to compute the sum

$$S = 1 + 2 + 3 + \cdots + 100,$$

and then sat back to enjoy a welcome break. To the teacher's surprise, Gauss returned in a few minutes claiming that the sum is 5050. Was he right? If so, how did he compute the sum so quickly?

Gauss noticed that the sum remains unchanged when the terms are added backward from highest to lowest. In other words,

$$S' = 100 + 99 + 98 + \cdots + 1$$

equals S. Then Gauss added the two sums:

$$\begin{array}{r} S = 1 + 2 + 3 + \cdots + 100 \\ +S = 100 + 99 + 98 + \cdots + 1 \\ \hline 2S = 101 + 101 + \cdots + 101. \end{array}$$

I should have thought of that... But it would have taken me a while, and I'm pretty sure I couldn't have done it at age 3.

This is pretty awesome - in fact, I've never seen this before. Outsmarted by a 3 year old...

Why was he going to school at the age of THREE? Did any one else besides me not know addition at the age of three?

I love how we all went wild over this. It's always nice to have the reading sprinkled with interesting anecdotes.

I've also never heard this anecdote before or even knew this trick, it's really interesting and I'm glad it's in the book! Even though it does make me feel like I pale in comparison to a child...

I remember learning this method years back, but I'm pretty sure it took a while for me to figure it out.

About being outsmarted by a 3-year old – that's my daily life these days. As other parents of toddlers will probably agree, any time I manage to win an argument with my daughter (who is 2.7), and my winning happens rarely, I cheer to myself, "Yeah, I outsmarted a 2 year old!"

Dude that's wicked smart. I would not know what to do if my kid was that much a brainiac.

I don't know if you have more examples of symmetry... but this section is short and this example is pretty awesome, maybe insert another example?

Additional examples would be nice, though this I think is suitably sufficient for its purposes. I feel like more examples would require additional content beyond just that, which is pretty time-intensive to do.

So I haven't done this problem in a while, and didn't remember the formula, but from the reading it's easy to see the formula is n(n+1)/2

I have always best understood this problem graphically with two triangles combining to be one rectangle, does anyone else look at it that way?

I do

3

Symmetry and conservation

3.1	Heat flow	46
3.2	Cube solitaire	48
3.3	Drag using conservation of energy	50
3.4	Cycling	54
3.5	Flight	56

Symmetry greatly simplifies any problem to which it applies – without any cost in accuracy. A classic example is the following story about the young Carl Friedrich Gauss. The story might be merely a legend, but it is so instructive that it ought to be true. One day when Gauss was 3 years old, the story goes, his schoolteacher wanted to occupy the students for a good while. He therefore asked them to compute the sum

$$S = 1 + 2 + 3 + \cdots + 100,$$

and then sat back to enjoy a welcome break. To the teacher's surprise, Gauss returned in a few minutes claiming that the sum is 5050. Was he right? If so, how did he compute the sum so quickly?

Gauss noticed that the sum remains unchanged when the terms are added backward from highest to lowest. In other words,

$$S' = 100 + 99 + 98 + \cdots + 1$$

equals S. Then Gauss added the two sums:

$$\begin{array}{r} S = 1 + 2 + 3 + \cdots + 100 \\ +S = 100 + 99 + 98 + \cdots + 1 \\ \hline 2S = 101 + 101 + \cdots + 101. \end{array}$$

o wow ingenious (especially for a 3-year-old boy)

Yeah i always new the formula on how to get the sum, but I never knew/always forgot how it was derived. Good example!

This is a wonderful first example of symmetry.

the way i've heard this story told is that he realized that if he took pairs that summed to 100 (0 and 100, 1 and 99, 2 and 98, 3 and 97, etc), he'd have 50 pairs of 100 and a 50 left over.

This is the way I would go about doing the problem but it is pretty much the same thing. Also it would be fifty pairs that add to 101 not 100.

This is a neat impressive trick that can be useful for summations . I wonder if the story is actually true..

This version is also correct - it's 100 pairs of 101 that add up to 2S.

50 pairs of 100 + the left-over 50 as mentioned above is also correct.

yeah that's what i heard too, but same difference.

I hope that we see an example of a more complicated calculation using this method. Doing it with a simple sum is pretty intuitive, but I think right now I'd be hard-pressed to use this for anything else.

In this form, 2S is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

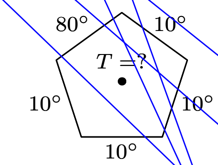

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa\nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa\nabla^2 T = 0$.

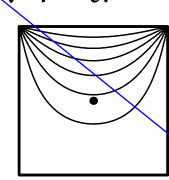

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at 10°, 10°, 10°, and 80° (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every 10°). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

Is this class taught in such a way that the previous concepts build upon the upcoming ones? It seems we used D/C to explain abstractions and are now using abstractions as a stepping stone for symmetry

> I have tried to design it that way, and am happy that it is showing results. One reason I put abstraction early on is that everything else in the class is an abstraction.
>
> In a couple chapters from now, we'll use symmetry to understand dimensional analysis. Then the limits of dimensional analysis will lead us to extreme-cases reasoning.

Maybe it would be useful to put the actual formula for the sum after we have been given the symmetry argument, to really show the power of representing symmetry into non trivial formulas

I find this example easy to understand.

that is a interesting way to do this, its the kinda thing that can be useful in many different types of situations

I've actually never heard of this trick before, thats awesome! But it only works for finite sets (obviously, I suppose) which means it's only applicable sometimes. Still very cool.

Well not exactly... I used this for computing the sum of 1 + r + r^2 ... more or less... sorta a combination of abstraction and symmetry. You can multiply this sum by r and you have the same sum - 1. If you add them together you will surely find your answer.

I remember learning about this in HS when we learned about series

This is really useful. I feel like we used this a lot in 2.001.

I prefer a similar trick that I find slightly more intuitive... There are 100 numbers, with average 101/2. So their sum is 100*101/2, which results in the same math, but is faster for me to apply in other occasions.

Before I read the explanation, I was reasoning it this way: 1+99=100, 2+98=100... all the way to 49+51=100, and 4900+100+50=5050.

I was imagining some kind of symmetry around the "50" axis.

In this form, $2S$ is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

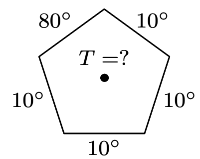

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa\nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa\nabla^2 T = 0$.

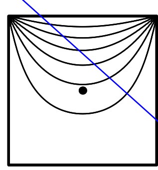

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at 10°, 10°, 10°, and 80° (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every 10°). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

Instructive, interesting example (though I'm not sure the story is that he was 3...it's very hard to believe).

For some reason I have a hard time realizing when to use symmetry. I suppose I should check that first every time, it's always pretty simple when it works.

He also says that it takes practice to get good at realizing when to use it.

I feel like symmetry and recursion are similar in this manner, the only wy to get good at them is to see a bunch of examples

Definitely. It seems more of a manner of thinking than anything out, just as in programming sometimes people understand recursive intuitively while others struggle to apply it - it depends on how you think through problems.

i feel like anything can be an abstraction at this point. any problem solving tool is an abstraction because it is separate from a specific problem? this makes abstraction rather....abstract. i think it's more useful to focus on the tools.

As you point out, all the problem-solving tools are themselves abstractions. Every word is an abstraction. Abstrations are everywhere! One purpose of this unit is to help everyone see how prevalent abstractions are and to start noticing them everywhere – as the first step toward using them and then making new ones.

At first when I read this I wondered why you didn't tell us how generally applicable abstraction really is. Now, I'm thinking we just wouldn't have understood without all of the examples.

I think of symmetry as a feature of the problem, not as the transformation that exploits that symmetry.

When an old physics teacher explained the idea of symmetry being a feature of the problem rather than something to try to force upon a problem it really cleared it up for me - I think its a good way to think about it that helps things make sense.

I like how symmetry is put into the category of abstraction here. I had previously thought of it as its own thing, but it really is just reusing information over lines of symmetry.

So what differentiates symmetry from abstraction? is it just repetition?

In this form, 2S is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

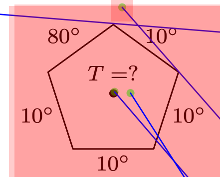

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa\nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa\nabla^2 T = 0$.

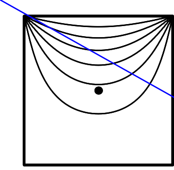

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at 10°, 10°, 10°, and 80° (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every 10°). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

I haven't ever dealt with heat flow problems before...not sure if this is a course 2 familiarity, but the symmetry approach makes it seem almost too easy. What a great example.

I totally agree with you...I finally get it!

I haven't ever dealt with heat flow problems before...not sure if this is a course 2 familiarity, but the symmetry approach makes it seem almost too easy. What a great example.

I agree- the examples in this section are excellent.

Heat flow implies that we might be calculating d/dt of something. Maybe a better title is "Temperature gradients" or something like that.

I agree. The title Heat Flow implies that there will be some transfer of heat, whereas you're really just interested in temperature differences and gradients.

If there's a temperature difference, there will be heat flow. The temperature distribution is just at equilibrium.

Oh, I just realized it's the temperature not the angles... Maybe that could be clarified as well. It is obvious when I read through the material but not when I glance at it.

This is a regular problem in 18.04...we use complex numbers symmetry together with regular physics equations

I am confused about how angles are notated here...

This is a regular pentagon, so each angle should have 360/5 degrees. The notation is actually the temperature at which the edges should be held at.

what are heat sources and sinks? A meche term?

In real life would there by anyway to keep the points between where the 80degrees and 10 degrees meet from changing each others temperature?

Could this even be easily approximated (divide and conquer style)? Or is it too complicated to easily break down like that?

I was wondering a similar thing - how do we decided to approximate something vs trying to find a clever way to find and exact answer?

In this form, $2S$ is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

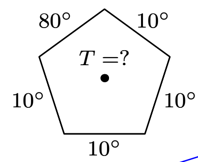

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa\nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa\nabla^2 T = 0$.

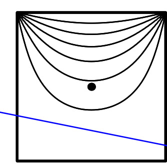

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at 10°, 10°, 10°, and 80° (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every 10°). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

can you explain what exactly brute-force is? I've also heard of this term, but I don't really know what it means.

Brute-force methods are those approaches where you know you'll get the right (usually exact) answer, but that will take more time/effort to complete. They usually involve either using a difficult analytical solution or some repetitive process that could be automated or collapsed into fewer steps.

The complexities involved in actually mapping the calculations and heat flow by brute force is nearly impossible. The equations would just be too complex.

The brute-force method is suggested below with the differential equation and explanation.

An explanation of what this is would be helpful.

The next sentence explains the variables and the sentence after explains the equation.

If I never took Thermodynamics, I would be a little confused here.

Is this constant specific to the type of material/thickness?

it's a material property

From wikipedia, thermal diffusivity is "the thermal conductivity divided by the volumetric heat capacity", so it's variable depending on the object.

This bit of information seems irrelevant to the question at hand. (Not that it's not a useful thing to know, just it's unclear why it's here.)

I think it's still relevant in that it explains why the dT/dt component of the equation drops out. If the text had initially introduced the heat flow equation in the time-independent/stable form then I would agree.

I think what Sanjoy wants to do is start with the original equation, and show how it simplifies to this easier-to-deal-with form where the right side =0.

redundant (you already used the word eventually in this sentence)

The exact meaning of this diagram seems a little unclear, and it isn't really referenced directly in the text. Maybe label it more clearly? And explain the contour lines?

In this form, 2S is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

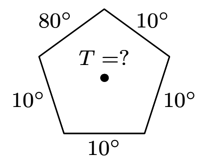

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa \nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa \nabla^2 T = 0$.

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at 10°, 10°, 10°, and 80° (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every 10°). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

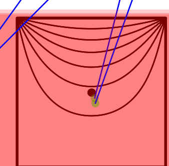

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

not the best example.

I disagree. I think the square is a good example because it is a simpler symmetrical shape that we can compare with.

I also think the drawing is good, but maybe a little more, like a directional, or some labels would be better. If you miss something in the reading, then you will completely skip the picture, and it has a profound effect.

I think labels here might help if someone were to quickly go back to look at it again. I always find figure captions very helpful for reviewing materials later.

i'm not really sure how this all makes sense

just refer to the diagram

I think it's fairly straightforward, since the figure is right next to it.

It sort of makes sense... there shouldn't be a gradient on the bottom of the diagram since the edges are the same, and you expect to see a larger gradient on top. I doubt I would have gotten the exact shape right, but seeing the shape, it makes sense

The idea behind the gradient makes sense, but I think the exact shape of the contours in the diagram is quite non-intuitive.

I find the shape to be somewhat intuitive, maybe not the spacing. I wouldn't imagine it being in any shape other then an arc.

I too agree that the diagram makes sense, but given the edge temperatures and asked to find the stable temp at a particular point, my answer would be very inaccurate. Therefore I would tend to agree that the distribution is highly non-intuitive.

I agree. It's intuitive that there is a gradient, but as far as the shape of it, I would have no idea. Plus, as the shape gets more sides, I'd expect that the gradient would get harder and harder to visualize

I would think that it turns out to be something like the above picture but with closely spaced lines towards the edges that are 10 degrees. the far 10 degree lines would have spacing farther apart than the ones next to the 80 degrees edge

In this form, 2S is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

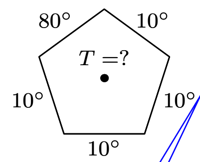

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa \nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa \nabla^2 T = 0$.

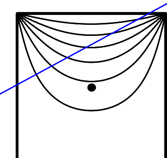

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at 10°, 10°, 10°, and 80° (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every 10°). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

Is this supposed to be a pun? If so, well done!

Is there a way to rotate something about the center and have the center change temperatures? I can't really think of anything.

Agreed - how would rotation every change the temperature at the center, even if the shape is not symmetric?

Well, if the shape weren't symmetric, the temperature at the center would change based on which heat sources were where. For instance, if one point of our pentagon were longer, then putting the hottest source there versus at a closer point would definitely change the answer.

wha?

blink what person?

I think this is not good word choice. "Nature, as exemplified by the heat equation..." or something similar to that

or "represented by"

I agree. "Person" is both distracting and confusing.

typo for persona?

This statement seems (for lack of a better word) obvious. I can't think of any systems in which the temperature would change in the center due to the directions of coordinate systems?

I am experiencing a bit of a disconnect. Even though nature doesn't care in what direction our coordinate system points, why do we have to consider the 5 directions that our pentagon can point. Why can't we just consider 1 orientation, add up the temperatures, and divide by 5?

In this form, $2S$ is easy to compute because it contains 100 copies of 101. So $2S = 100 \times 101$, and $S = 50 \times 101 = 5050$.

Gauss tremendously simplified the problem by finding a symmetry: a transformation that preserved essential features of the problem. The idea of symmetry is an abstraction, and fluency in its use comes with practice.

3.1 Heat flow

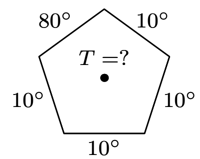

As the first example, imagine a uniform metal sheet, perhaps aluminum foil, cut into the shape of a regular pentagon. Attach heat sources and sinks to the edges in order to hold the edges at the temperatures marked on the figure. After waiting long enough, the temperature distribution in the pentagon stops changing ('comes to equilibrium'). Once the temperature equilibrates, what is the temperature at the center of the pentagon?

A brute-force analytic solution is difficult. Heat flow is described by the following second-order partial differential equation:

$$\kappa \nabla^2 T = \frac{\partial T}{\partial t},$$

where T is the temperature as a function of position and time, and κ is a constant known as the thermal diffusivity. Eventually the time derivatives approach zero (the temperature eventually settles down), so the right side eventually becomes zero. The equation then simplifies to $\kappa \nabla^2 T = 0$.

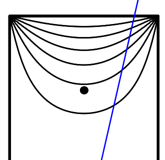

Alas, even this simpler time-independent equation has easy solutions only for a few simple boundaries. Even these solutions do not seem that simple. For example, on a square sheet with edges held at 10°, 10°, 10°, and 80° (the north edge), the temperature distribution is highly nonintuitive (the figure shows contour lines spaced every 10°). For a pentagon, even for a regular pentagon, the full temperature distribution is still less intuitive.

Symmetry, however, makes the solution flow: Rotating the pentagon about its center does not change the temperature at the center. Nature, in the person of the heat equation, does not care in what direction our coordinate system points. Stated mathematically, the Laplacian operator

What exactly is the definition of this symbol?

This operator is from 18.03 I believe- from Wikipedia- The Laplace operator is a second order differential operator in the n-dimensional Euclidean space, defined as the divergence of the gradient.

The Laplace operator is div(grad).

It took me a little while to make the connection between "Laplacian operator" and Laplace...maybe I just need a little more sleep...

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

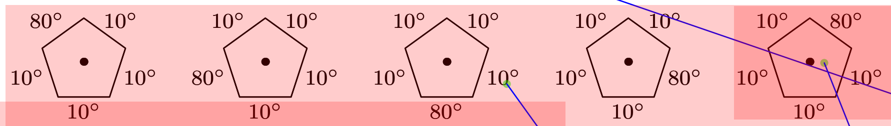

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{edge} = 80^\circ + 10^\circ + 10^\circ + 10^\circ + 10^\circ = 120^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120°, so the temperature throughout the sheet is 120°.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72°. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

can't you just say it doesn't have an inherent coordinate system rather than making us think about spinning pentagons?

I think you need the pentagons in there to show what the symmetries are. A scalene triangle has no symmetries but the Laplacian is still independent of the coordinate system.

Other than doing this rotation part, couldn't you have simply "weighted it"? Known that the temperature was going to have 4 parts influence from 10 and 1 influence from 80, and added to 120 and divided by 5 for this reason?

this is what i would have naturally thought to do as well. is this a correct line of logic?

I think so – both methods seem to be using the same idea, but I think the one expressed in the text is more intuitive. Your mileage may vary, though.

The idea that the temperature moves in the way the pictures show seems pretty intuitive to me. It would be nice to see a comparison between this and the way Gauss was able to flip the numbers backwards then add them up.

It's the same concept though, averaging over different perturbations of the same problem to figure out what one answer would be from multiple solutions

To me, it is clear that the pentagon is rotationally symmetric, more clear than the operator being invariant. Perhaps you could say "By inspection, the pentagon is rotationally equivalent. This fits with the definition of the operator, which is also defined as such"

I think the visual presentation here is good. I remember looking at the tree drawings and thinking they look messy and drawn in paint, but this looks very professional

But they were drawn with his "tree" language!

it would be great to have a diagram with the line spacing to illustrate the temperature distribution.

This example was explained VERY well and made perfect sense. It was clear, concise, and easy to understand.

I dont understand this part

I wish I had known this trick when I took 18.303. We had to do things like this the long way...

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

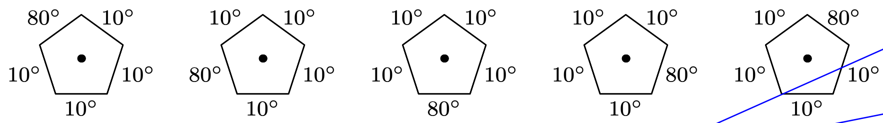

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ = 120\,^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at $120\,^\circ$, so the temperature throughout the sheet is $120\,^\circ$.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by $72\,^\circ$. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

I would have just guessed that since the dot is in the middle, that each edge contributes equally. I forgot if you said it in class, but is this any different from stacking them?

How do we know that temperature addition obeys linearity? I would have approached the problem by averaging the temperatures of the edges, which I know works – e.g. mixing hot and cold water – and that does give the same answer.

On second thought, since the differential equation is linear, I suppose both approaches should be the same, though I still don't like adding temperatures because that is not as rooted in a physical phenomenon.

Here's is where an abstraction is useful. Blur your vision a bit and don't look too closely at what you are adding. The method works because of a general property of the differential equation (linearity), independently of what the particular quantity being differentiated means.

Alternatively, think of *averaging* the temperatures rather than adding them. Averaging temperatures is a physically legitimate operation.

I think it is fairly obvious to me that it would obey linearity, just from experience. Also, averaging is a linear operation, so aren't you also assuming temperatures add linearly?

This still seems counter-intuitive. If I have a 2 sheets of metal at 100deg and stick them together, their temperature doesn't become 200deg...

More explanation regarding how this is legitimate would be fantastic...

I agree with this comment; it does seem very counterintuitive since temperatures don't magically add when you stack them together.

I am also very confused as to why using symmetry enabled us to stack the pentagons as a way of simplifications..

I thought when materials of different temperatures come together, they would reach a certain equilibrium. I don't understand why we can add up all teh temperatures

I'm just struggling to conceptually understand why we can add all the rotated pentagons to come to the right answer. Maybe I just don't understand this Guassian principle.

I think this brings up a good question. In general, when do we know that we can assume a problem to be symmetrical?

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

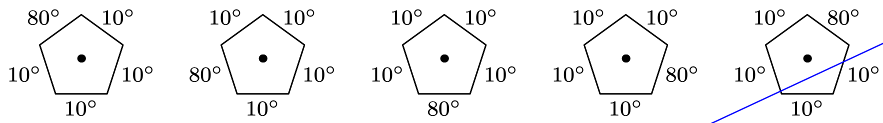

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{edge} = 80^\circ + 10^\circ + 10^\circ + 10^\circ + 10^\circ = 120^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120°, so the temperature throughout the sheet is 120°.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72°. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

i'm not following this logic

temperature is summable???

This is really clever, actually. But I don't think I'd be able (even after seeing this example) to be able to apply this idea–can we use it any time when we have things evenly spaced from a certain point? and in this example–even if the points weren't evenly spaced, could we still use the same method? how would we take the average for the temperature "at the edge" if the edges were different distances from the center?

You couldn't. This only works (I think) because of symmetry. Once the object in question is no longer symmetric, this trick no longer applies

Is it possible to pick a different point to find the temperature at using symmetry? Or is it generally always at the center?

Can you find another point where the figure is symmetric?

You can make any number of arguments about two points having the same temperature, but without the full symmetry of the central point, you can't figure out easily what that temperature is...

This is so cool! I love seeing Gauss's method applied in a way I never would have thought of.

gauss' method? not really...it's symmetry which i'm sure was not invented by him.

This is very clever, clearly written, and unites the idea being presented here to the introduction about Gauss.

I thought the edges are also stacked on top of each other? maybe I am thinking of a wrong picture

They do..all of the edges sum up to 120deg.

I am so confused as to how this is able to work- I understand the math, but temperature??

This wins the award for my favorite problem we've done so far. Definitely elicited the biggest "I didn't see that coming" reaction yet.

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

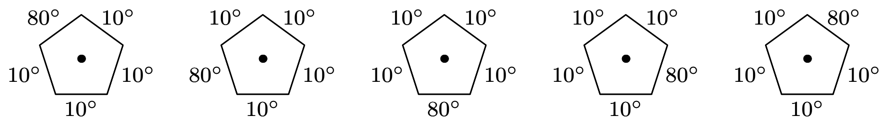

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{edge} = 80^\circ + 10^\circ + 10^\circ + 10^\circ + 10^\circ = 120^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120°, so the temperature throughout the sheet is 120°.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72°. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

Can you just add temperatures like that?

I was wondering the same thing...at one level it sounds like it should be good but it still seems too easy to be true.

I agree especially when the problem says the edges are held constant with sinks and sources. Wouldn't the sinks just eat up the 80 degrees?

I guessed this would be the solution. But it just seems so simple that I didn't want to trust the answer.

Well, didn't the professor say that part of this class was learning to trust our gut?

This only works because we're talking about the system at equilibrium, right?

okay, so it makes sense intuitively to divide by 5 – but what I don't understand is if the temperature is 120 throughout, why does having the 5 centers aligning mean you need to divide by 5, shouldn't it just be the same throughout?

The way I see it is if you stack the 5 pentagons on top of each other then the sums of the temperatures at a point on the combined pentagon will be 120 degrees. Since the center of each pentagon has the same temperature (it doesn't change with the orientation of the heat sources) the center of each pentagon must be at a temperature of 1/5 of 120 degrees.

Hope that helps.

That does make sense, but I thought the temperature distribution was uniform and at equilibrium?

noooo the temperature is no longer CHANGING but that does not mean its uniform. (See the above temperature gradient plot for the square). What this says, is that the ONLY point on these 5 overlaid pentagons where each perturbation would have exactly the same temperature is in the middle, where we're looking, so you can in fact average every possible rotation and divide by the number of rotations. It would not work if it was not the center, or if we had assumed the temperature gradient was still changing

all right, that makes more sense.....the reason why I assumed it was uniform was because in the previous paragraph it says "so the temperature throughout the sheet is 120"...that should be changed so it is not so confusing.

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

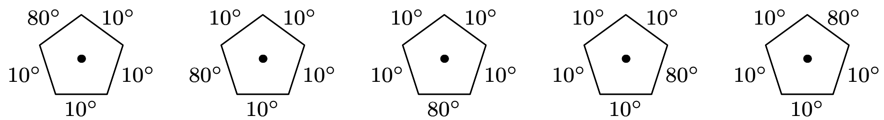

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{edge} = 80\,° + 10\,° + 10\,° + 10\,° + 10\,° = 120\,°.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120 °, so the temperature throughout the sheet is 120 °.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72 °. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

I would have guessed that even though I don't know much about thermo

so this is a simple add and distribute heat problem since each side is equidistant from the center.

this is the kind of stuff I was hoping to learn in this class. I hope I can start thinking like this.

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

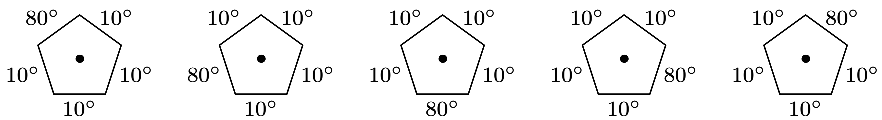

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80° + 10° + 10° + 10° + 10° = 120°.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120°, so the temperature throughout the sheet is 120°.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72°. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

I'm not at all convinced that this is the correct answer. Honestly, it felt like we were arbitrarily stacking pentagons and adding their temperatures for no particular reason. I understand that temperature doesn't change under rotation, but how does that imply that we can stack them and add their temperatures?

Is this the same as just finding the average of the pentagon's sides temperatures?

yeah that would've been an easier explanation

Agreed...averaging the temperatures of the edges is the same argument, and much easier to follow. I think he was just looking for a symmetry example that was easy to grasp, but perhaps this is not the best one...

I agree as well - this explanation had me thinking "Wow! _IF_ that is true, that's a cool method of solving it" with a big emphasis on the if. I wasn't convinced as to its validity.

Um, given that it was already stated the temperature at each side is set (forever) at 10 or 80 degrees, this would not work to find the temperature at a side (that is already known). it means that instead of solving complicated equations and looking at gradients...at ONE point (the symmetry point, the center) we can average over 5 possible configurations which would all give the same center value. NOW (since all 5 have been averaged) the whole pentagon is at one temperature, including the center. So then you divide out by the number of times you had to rotate it to get back to the center temp for ONE pentagon and have the answer

It ends in the same calculation, but the symmetry explains why it works to average the side temperatures. The symmetry argument also shows what the necessary ingredient is: a linear differential equation. If heat flow had been described by a nonlinear differential equation, then it would still seem plausible to average the side temperatures but it would no longer be correct.

I think that one point that might help is that the edges are heat sources. If you put all the sources at one point (or edge) then the resulting temperature output will be the sum of the individual temperatures.

I'm not sure this holds here.... I think this just gives the average temperature...not necessarily the temperature in the the center of a particular pentagon

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

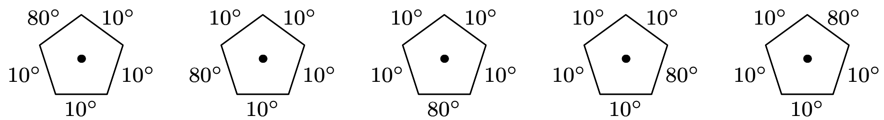

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{edge} = 80\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ = 120\,^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at $120\,^\circ$, so the temperature throughout the sheet is $120\,^\circ$.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120^\circ/5 = 24^\circ$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by $72\,^\circ$. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

I need a little more help with how the heat transfer stuff is the same as the gauss's sum

This symmetry method seems very useful, but I feel like it would be extremely hard to intuitively devise- or find- a symmetric way to solve a problem, especially a hard problem.

I agree... I would have never thought of this approach... I think with more practice with symmetry problems this will become more intuitive....hopefully.

Yeah I agree...I think it will become a lot more clear how to use this method if we go over some more simpler examples in class.

I agree. I understand it now that it's worked out, but I'm not sure I could think to do it this way given another problem.

but you didn't actually solve for the temperature curves, so you wouldn't have used that anyway. and your statement that it was "non-intuitive" doesn't apply when you don't solve for the curves either. we only looked at one extremely convenient point.

It's more that he took the only possible point. This wouldn't work for elsewhere in the pentagon, but for this point (symmetric) it does work, and that's the beauty of it. The Gauss example is clearer, but only applies to a very limited number of cases, whereas this could be applied to many, but only to solve for one point.

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

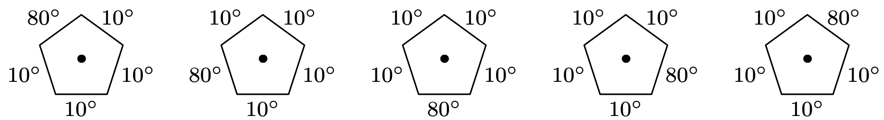

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ = 120\,^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at $120\,^\circ$, so the temperature throughout the sheet is $120\,^\circ$.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by $72\,^\circ$. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

I find it interesting that both problems are actually solvable in other ways. Is this always the case? Is there ever a case where using symmetry or some derivative of symmetry is the ONLY way to solve a problem?

I doubt it, but I do feel that symmetry makes things a heck of a lot easier.

I agree. I can't think of a situation in which symmetry allows you to solve a problem that is otherwise impossible (assuming infinite computational resources, of course).

This may not actually answer your question, but symmetry is extremely important in mineralogy and crystallography. The space groups and all those groups related were at least originally solved used symmetry - and mineral structures can also be solved using symmetry. I'm not really certain of more modern, computational methods, so there may be ways to do it without symmetry now - but I just don't know.

I don't think so. I think it's just a shortcut for problems that are normally more complex.

I agree, just as in Gauss's problem, it is possible to solve the problem with computational methods but we are learning quick computations.

I'm sure there are always multiple ways to solve a problem, but it looks like whenever symmetry is possible, it is the easiest–the trick is just being able to identify the symmetry

what, precisely, is a "symmetry operation" in most concise words? the act of.....?

only after some time was I able to understand the solution. Its not immediately apparent

This paragraph is very good. It sums up the lessons from these two problems. Sometimes these examples can get caught up in the details of the physics/engineering that the approximation lessons is lost. These types of paragraphs are great.

not necessary.

"rotated on each of its sides" would fit more with your goal of making it appear simple.

I disagree. It's precise than saying "one rotation" or something like that.

*more precise

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

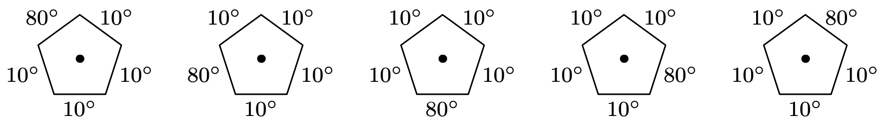

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80^\circ + 10^\circ + 10^\circ + 10^\circ + 10^\circ = 120^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120°, so the temperature throughout the sheet is 120°.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120^\circ/5 = 24^\circ$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72°. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

I really like this explanation and how it ties back to the Gauss story - flipping the pentagon is just like flipping the number sequence - kind of gives new meaning to lines of symmetry for me.

I would like to see more examples. I feel like symmetry is best grasped through lots of practice and lots of exposure to different ways of applying symmetry arguments to different types of problems.

This is analogous to controls in a scientific experiment.

Possibly putting this before the examples would be helpful

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

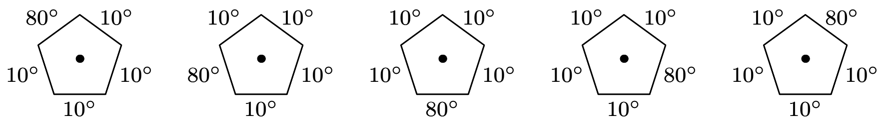

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ = 120\,^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at $120\,^\circ$, so the temperature throughout the sheet is $120\,^\circ$.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by $72\,^\circ$. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

This is a good summary and it answers my question from earlier. Are you going to add practice problems to this section like you did in previous sections? I think they would be very helpful since this is a concept that requires a lot of practice to master.

While this is a good summary of the section, it would have been nice to see the example problem broken down in to a summary like form, in order to have more clear insights as to the thought process of using symmetry to solve a problem.

I'm a little confused here - what exactly "did not change" in this example? The temperature at the center when we rotated the shape? I'm also still failing to see how symmetry makes the problem easier when we could have just averaged the temperatures of the edges (although that also uses a symmetry argument - perhaps use that instead?).

Perhaps this concept of invariants should be a section in and of itself? It would greatly simplify certain classes of problems.

I agree that the idea of the invariant can be discussed further. I hope to see it expanded in other sections, or else I suggest the concept of the invariant be expanded. It makes sense, but I'd like to see more examples.

I think that the part that didn't change was that the Laplacian operator doesn't change with pentagon orientation so by adding multiple pentagons to each other all sides are now 120* which makes the problem simple. Same with the Gauss problem, adding the numbers doesn't change if you add them up or down, so you can add both ways together and divide by 2.

I really like how you gave a strategy for examining problems, and think it is a great way to wrap up this section and go into more detail about symmetry.

I agree, previously we would often have examples explaining a particular concept but it was sometimes difficult to extract the actual lesson from the example. I think this summary does a good job of reinforcing the concept and also leaving the reader with a practical application point that one can remember.

What I am learning from this thread (and many others) is that it is very helpful to reap the fruits of each example. In other words, after doing an example, figure out what are the transferable lessons (and maybe do another example with those lessons explicitly). I'll try to do that more as I revise the book.

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

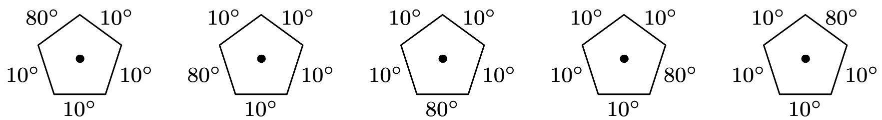

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ = 120\,^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at $120\,^\circ$, so the temperature throughout the sheet is $120\,^\circ$.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by $72\,^\circ$. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

I'm really excited to see how this book is going to end up!

This sentence cleared it up for me. But I don't like how I have to wait until the end of the article to get to the moral of the section. The path that you take us through could be much simpler.

Likewise...the mention of invariants brought back other examples of invariants that have been used for problem solving and opened up a lot of other ideas of symmetry applied - it'd be nice if this was brought up earlier so that we could have other applications in mind too

Actually, I'm not sure this statement would have made sense to me if I hadn't seen the example first.

I agree, it's easy to understand what he's talking about after seeing the example, but I don't know if would be this clear in the beginning.

I do understand this summary here, but what are some other examples of problems where symmetry can be used? and what is unchanged/changed? it looks like addition problems and anything involving shapes (as long as there is constant distance from a certain point/center) can be solved with symmetry, but what else?

Additionally, I feel like if I had seen this problem I wouldn't have tried solving it symmetrically this way because I would have convinced myself its more complicated than that. How can you tell when it really is that simple? (and when something looks like you can use symmetry but shouldn't)

Good Summary!

Will we always solve for these unchanged quantities by adding up the symmetrical results and dividing by the total number of symmetries (as we did in both of these examples)?

I feel like that there's some abstraction here... so probably not. I guess I think the point more is to look for symmetry and see how you can use it in general, not that there is always this specific formula for symmetrical things.

These examples were both easy to follow.

Try not to highlight the entire paragraph with a comment box. It makes it difficult to isolate the comment boxes underneath it.

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

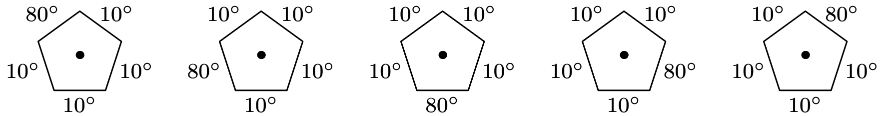

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80^\circ + 10^\circ + 10^\circ + 10^\circ + 10^\circ = 120^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at 120°, so the temperature throughout the sheet is 120°.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by 72°. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

why won't these stupid boxes stop popping up???

Which boxes? It sounds like an NB bug. Can you send me and Sacha <sacha@MIT.EDU> an email describing specifically what happens?

Me 2. when I'm writing a reply and I click with the mouse. Another "options" box pops up.

I've found the way to avoid this is to use the arrows to move around while you're typing instead of clicking. If the box pops up, you can click in another part of your answer to move it around so you can see your text or click the Save button.

Yeah this happens to me all the time too...I think its a bug in the NB user interface.

3:47 ... yes, that is a way to get around it, but it's still annoying ... and I'd really like to be able to highlight something I'm writing easier.

to be more specific, this box pops up when we're typing a reply and then click inside the box to, perhaps, go to a different part of our text to readjust wording. the cursor still goes back but another box pops up to give us the option to reply to it. quite annoying.

Ah, that must be it. The 'reply' dialog box somehow does not shadow the comments underneath it. So a click in the reply box is interpreted by NB as a click on the comment underneath it, so NB then offers another reply box...

I never noticed it because I never use the mouse unless I have to, so I always used emacs editing keys (ctrl-P to move up, ctrl-N to move down, etc.) to edit comments.

But I agree, it is an NB interface bug. I suspect it is easy to fix. I'll point Sacha to this thread.

Hi everyone, Thanks for the feedback... I'll try to fix that ASAP.

OK... so it looks like it's related to the "options" menu: While this gets fixed, a workaround is to bring the context menu not by clicking on the "options" link but by RIGHT-CLICKING anywhere on the note (mac users: read control+click)...

∇^2 is rotation invariant. Therefore, the following five orientations of the pentagon produce the identical temperature at the center:

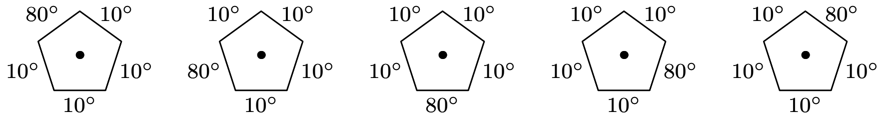

Now stack these sheets mentally, adding the temperatures that lie on top of each other to make the temperature profile of a new metal supersheet. On this new sheet, each edge has temperature

$$T_{\text{edge}} = 80\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ + 10\,^\circ = 120\,^\circ.$$

Solving for the resulting temperature distribution does not require solving the heat equation. All the edges are held at $120\,^\circ$, so the temperature throughout the sheet is $120\,^\circ$.

That result, with one more step, solves the original problem. The symmetry operation is a rotation about the center of the pentagon, so the centers overlap when the plates are stacked atop one another. Because the stacked plate has a temperature of 120° throughout, and the centers of the five stacked sheets align, each center is at $T = 120°/5 = 24°$.

Compare the symmetry solutions to Gauss's sum and to this temperature problem. The comparison will extract the transferable ideas (the useful abstractions). First, both problems look complex upon first glance. Gauss's sum has many terms, all different; the pentagon problem seems to require solving a difficult partial differential equation. Second, both problems contain a symmetry operation. In Gauss's sum, the symmetry operation reversed the order of the terms; in the pentagon problem, the symmetry operation rotated the pentagon by $72\,^\circ$. Third, the symmetry operation leaves an important quantity unchanged: the sum S for Gauss's problem or the central temperature for the pentagon problem.

The moral of these two examples is as follows: *When there is change, look for what does not change.* That is, look for invariants. Then look for symmetries: operations that leave these quantities unchanged.

It should be fixed now. Be sure to reload the page using Shift+Reload to that Firefox doesn't use its cached copy of the code, and give it a try

Thanks for letting me know, and have a good w-e.

I like the diagrams in this reading, the previous ones, and the next one (went ahead). They are very helpful and very simple; it makes understanding the readings easier, and also gives some nice supplement.

3.2 Cube solitaire

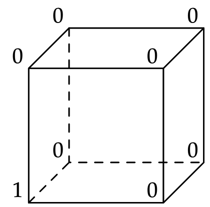

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

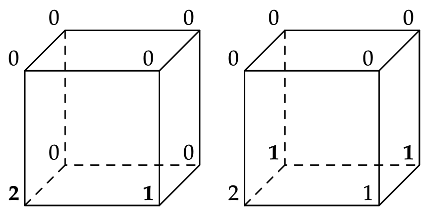

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

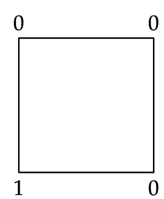

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

this definition of an invariant seems different from the one in previous sections?

The explanation above definitely helped understand how each variable(side) was configured to equal one. I think we should further explain the equation in class, step by step.

I think this example is great since it draws on things we all learned in 8.01. However, I am not really sure of its connection to the cube problems. I am not sure of the comment "physic problems are also solitarie games".

3.2 Cube solitaire

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

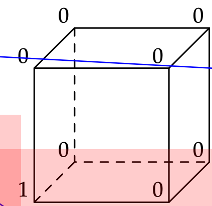

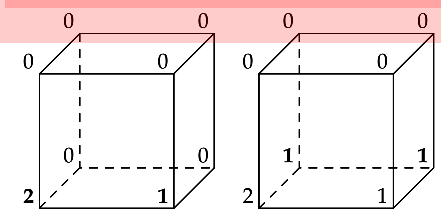

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

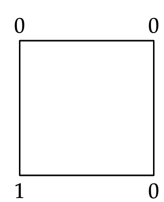

Read Section 3.2 for the memo due Sunday at 10pm. (Note: NO lecture on Monday.)

When I think of solitaire, I think of the card game. Is this some other meaning? Or is the term more general?

I'm also very confused here - wikipedia-ing solitaire comes up with the card game as well. What does this have to do with the cube?

I'm assuming it just means it's a game that you can play by yourself....

Ya, in this sense solitaire means a one person game/challenge. You guys never played solitaire games like this (http://www.woodtoysonline.co.uk/Family%20games_files/s where you had to remove all but the last peg by jumping pegs over others to remove them? I wonder what the invariant is for those games?

As in, in the last move, all the vertices must because 3? I think I'm misunderstanding the rules...

Nevermind, didn't see the "multiples" of part.

If this really is a game called solitare, is there a way nevertheless to call it something else to avoid confusion?

a followup question to this section could be for which n's does the problem hold?

This question is great example for this section. I originally tried to do it before breaking it up into pieces. But once I saw the problem in a different light, it was obvious.

this seems really similar to the checkers board covered in dominoes. I don't think you can win this game

I still don't really get the game

I wouldn't consider this a game...even a solitary game. I think that it would be better to call it a Puzzle.

3.2 Cube solitaire

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

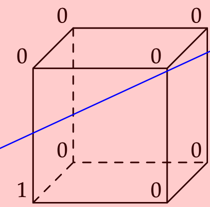

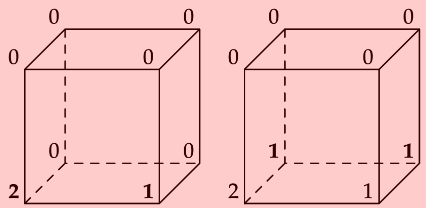

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

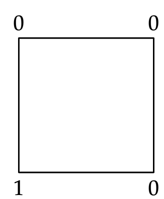

it might be nice to have animations for this so that a person could play the game and just see how it works.

or maybe one edge highlighted by a color or bolded and then the vertices circled. i know i understood the problem without any bold/circling but it's nice to know i understood it perfectly

I agree, although it's just as easy for me to take notes physically which is what I'm doing to keep it all straight.

I like the bolded suggestion if this was to be in a textbook instead of simply online.

Bolding the line would be useful, but it's not too difficult to follow and suggesting "animations" just seems ridiculous. I mean, you chose a line and its two vertices increased in value by 1. It's pretty straight-forward.

I think that the statement was more of a 'this would be cool' than a 'i need more to understand this" ... I, personally, would enjoy playing around with an app [just for "fun"]

ps. to orig. poster, please don't highlight huge blocks of text like that...it makes it harder to pull up notes for the one lines underneath it...do something similar to what you see along the edges. Thank you.

I agree, I had to read this over several times to understand it. perhaps the above suggestion would have helped.

I feel that if you have the technology why not make animations that can easily explain a point.

I showed this to a friend who likes puzzles, he looked at it for about two minutes and then told me it was impossible, and a used a variant on this proof. (I feel kind of dumb in comparison now)

So how many moves do we get? As many as we want?

I think we get as many moves as we want, but the phrase "neither configuration wins the game" doesn't make it super clear that from these positions and after many more moves, the game can't be won.

3.2 Cube solitaire

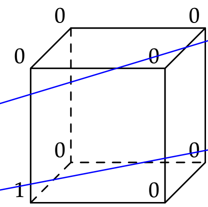

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

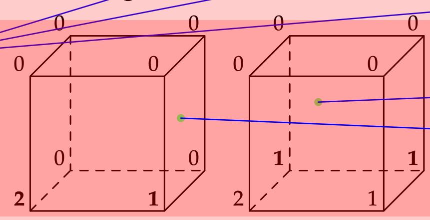

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

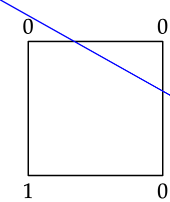

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

Not that it actually matters since it's impossible, but you should make it clear that any number of moves are allowed. It almost seems like here you indicate you can only have two.

Not that it actually matters since it's impossible, but you should make it clear that any number of moves are allowed. It almost seems like here you indicate you can only have two.

Not that it actually matters since it's impossible, but you should make it clear that any number of moves are allowed. It almost seems like here you indicate you can only have two.

I don't know about for others, but this is still not that clear to me... I think animation might actually help for people who are confused and are slow at understanding things.

I'm still confused about how this example as a whole illustrates symmetry. Any thoughts?

The first thing that came to mind for me was actually flattening the structure into a series of squares.

As for the image this comment is highlighting... I think it works well for understand how the game works. Most people seem pretty confused about the name "Solitaire", but after describing how the game works, and confirming the mechanics with this image, its very clear. I don't think further clarifications are needed to explain the game.

I understand how the both this game and the game of solitare are played but I don't really get why this game is called solitare.

I also think that this might not be the most effective way to begin explaining symmetry. The reading gets very dense when discussing mod and also trying to remember symmetry, and remember the solitare game. Its gets difficult trying to process, understand, and remember all of these while at the same time trying to make connection between them (i.e.. connection of how the game relates to symmetry). I just compared this to the next reading and for some reason it is much easier to read and comprehend than this. Just wanted to give some constructive feedback

Not that it actually matters since it's impossible, but you should make it clear that any number of moves are allowed. It almost seems like here you indicate you can only have two.

3.2 Cube solitaire

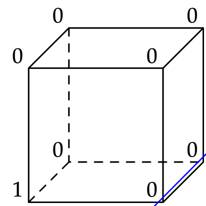

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

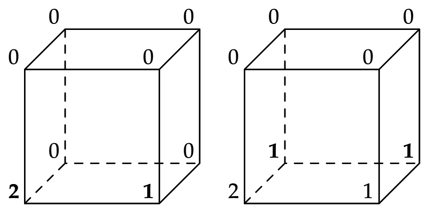

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

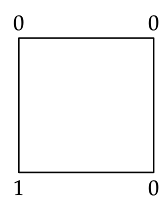

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

Try to answer this question for yourself before reading onward!

This might seem like a silly request, but it's actually useful for learning. Oftentimes we glaze over things with a "Yeah, I can do that if I wanted" attitude at the expense of actually doing it and learning. Active reading = good reading.

One advantage of personal teaching, including lecturing, compared to books is that you can stop and make sure everyone tries the problem. As the next-best thing, I'll put a solitaire problem on the next problem set. (We'd do one in class if I were not away on Monday.)

Can we do one in class Wednesday anyway?

I wonder if you could include one more example where this principle of the invariant helps you prove that a final state is true, as I'm curious as to how it would resolve into a winning sequence of moves?

I accidently read someone else's comment about this not being solvable before this. Perhaps the creators of NB could make it so that the instructor's highlighted boxes could be a different color?

Yeah. Your email address shows up highlighted in red, but not until the comment is expanded, and they all start folded in so we can't tell which ones are yours without flipping through them all.

I've been trying to work out a solution but it seems hopeless! Can we go over this example or another in class and prove that this doesn't work (as I assume)?

I think, with these kinds of questions, you can usually guess that you can't win. It makes the problem more interesting, and it's fun watching people hopelessly try to find solutions.

The way I see it, is that you always have to increment two points, and there are an even number of vertices on a cube. Therefore it's impossible to get an even number of total increments (sum of the number on each of the vertices = even) when you start out with an odd count (1) and only increment with even numbers (1->3->5 etc)

Looking at this problem again it's amazing to see how now we have methods to actually attack it. When I saw this weeks ago I spent too long playing with the 3D cube picking random sides and getting nowhere, but now I have a starting place

3.2 Cube solitaire

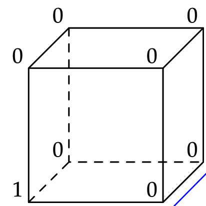

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

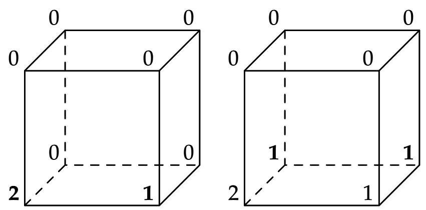

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

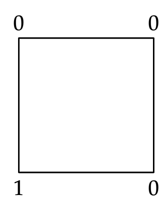

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

haha. i find it funny that this was the next sentence because i stopped reading just before this sentence, to try the game, as suggested, and stopped not too long later b/c i was overwhelmed by the options possible.

haha I did the exact same thing!!!

At first glance this sounds like a quick twist to a 6.006 problem

Could we maybe code this up and run it on the computer?

how did you figure this out?

ultimately you need the corners to sum to 24 unfortunately no multiple of 2 is within one of 24 (and divisible). if the sides were going to be 9 corners then it might be possible because you could at 26 points with moves and use the initial one value corner to get a total of 27 which is divisible by 3

Shouldn't this say "there are 12^10 sequences for ten moves." The current phrasing makes it sound like there are only ten moves possible in the game.

Or even a MUCH lower number!

Agreed! Anything somewhat lower than 12^10 is still huge.

how low can you go?!

I think it would help here to discuss the difference between combinations and permutations briefly.

This seems unnecessary

i suppose this depends on whether or not you're solving something else haha

I use this approach a lot when trying to figure out combinations and permutations, just write them out systematically in all possible ways, but it gets unreasonable if there are more than 20 or 30.

3.2 Cube solitaire

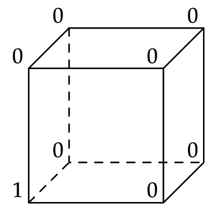

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

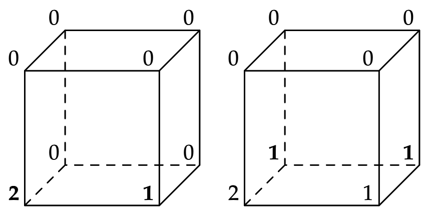

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

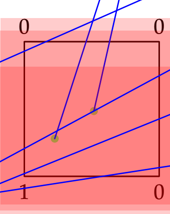

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

Simplifying the problems is always a great approach to finding a solution

Although this is true, it's always easier to see and understand the simplification rather than come up with it yourself.

I agree. I don't think I would have been comfortable reducing the 3D case to the 2D/1D case without being shown it wouldn't change the outcome.

Does that answer mean we can just simplify it and multiply our answer by the appopriate number?

i tried this! :)

yeah me too. I like how this problem goes back to the divided and conquer methodology.

I agree–perfect example of divide and conquer, and something that we could all do without a bunch of previous knowledge

how does solving for a square give us the answer for a cube?

Well, we can abstract with a decent degree of certainty that if there's no solution for a square, adding complexity will not make a solution _more_ likely. It doesn't give us an answer per se, but it lets us make some justified assumptions.

I like this idea of breaking the problem down into a smaller piece. I don't really think it's lowering the standard but just figuring out a strategy to apply to a simpler problem so you can apply it to the more difficult one.

SO does one just guess that trying this out with a square will lead to valuable information?

It could also be seen as a divide and conquer method

Makes sense

Makes sense

So I have some worries about simplifying the problem, and it's that I might skip over the key to answering the question. Ever since coming to MIT, problems aren't as straightforward. So I can't see myself really being able to use this outside of the class.

3.2 Cube solitaire

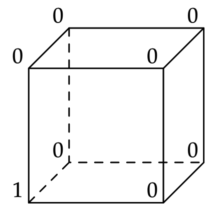

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

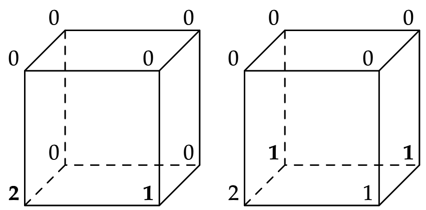

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

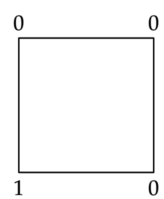

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

so this doesn't count as divide and conquer?

I think this might relate to the lossy/lossless situation that Sanjoy mentioned in class in reference to something else. Lowering your standards to me means simplifying the entire problem, while divide and conquer is breaking it into smaller pieces. But I'm making this up... I could be completely wrong. I guess they both relate in that you are taking something complex and making it simpler?

Yeah, I think it has to do with both–we are lowering our standards in that if we solve this problem, we've only solved the square, not the cube–a much easier problem. BUT,it could also be seen as divide and conquer since you're breaking the cube up into its faces. So if you could find the way that each face relates to the whole, you've solved the problem using divide and conquer. that's my understanding...

great statement

author of this quote? or is this one of your own?

I think this is one of the best lines so far in this class, though I feel like it's been stated once already. I think (and I think I've stated this before) that there should be a bunch of these; sure you lose some comedic effect, but it's memorable.

A proverb to live by! It seems like all the previous units kind of build on each other. This statement kind of reflects the divide and conquer approach as well. Instead of dealing with all 12 edges, divide it into just 4 and conquer that first.

That's hilarious.

...and this is why I love the readings for this class.

I love it.

I have a friend who is known for lowering his standards when the going gets tough!

I don't really agree with the statement, there is a difference between lowering standards and simplifying things. It does sound kind of catchy, and does sort of get the point across, but not the words I would have used.

This might be good for approximating, but I don't think it's a good proverb to live by. When the going gets tough, the tough should try harder! And perhaps collaborate/ask for help.

3.2 Cube solitaire

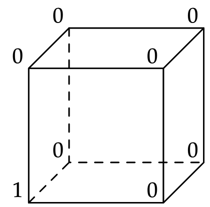

Here is a game of solitaire that illustrates the theme of this chapter. The following cube starts in the configuration in the margin; the goal is to make all vertices be multiples of three simultaneously. The moves are all of the same form: Pick any edge and increment its two vertices by one. For example, if I pick the bottom edge of the front face, then the bottom edge of the back face, the configuration becomes the first one in this series, then the second one:

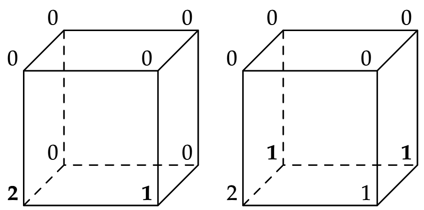

Alas, neither configuration wins the game.

Can I win the cube game? If I can win, what is a sequence of moves ends in all vertices being multiples of 3? If I cannot win, how can that negative result be proved?

Brute force – trying lots of possibilities – looks overwhelming. Each move requires choosing one of 12 edges, so there are 12^{10} sequences of ten moves. Although that number is an overestimate, because the order of the moves does not affect the final state, even a somewhat lower number would still be overwhelming. I could push this line of reasoning by figuring out how many possibilities there are, and how to list and check them if the number is not too large. But that approach is specific to this problem and unlikely to generalize to other problems.

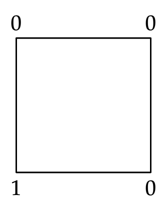

Instead of that specific approach, make the generic observation that this problem is difficult because each move offers many choices. The problem would be simpler with fewer edges: for example, if the cube were a square. Can this square be turned into one where the four vertices are multiples of 3? This problem is not the original problem, but solving it might teach me enough to solve the cube. This hope motivates the following advice: *When the going gets tough, the tough lower their standards.*

Are you really lowering your standards? To me, you're reducing the problem into something simpler upon which you can build to solve the original problem. I would say "When the going gets tough, the tough find a new going" instead of what's written.

We are actually relaxing the constraints of the original problem. So while your recommendation makes sense, the original is just as valid.

I dunno, I actually like the original quote as well. It fits the section a lot better.

Hmm this is actually very clever - you can represent all the vertices of a cube by two connected squares, so if you can solve a square, you can solve a cube (by never choosing the vertical edges).

I think he was trying to add an element of comedy into the text.

yea who cares if it's not 100% accurate of the precise situation. we get it, it's funny, and many of us like it.

The square is easier to analyze than is the cube, but standards can be lowered farther by analyzing the one-dimensional analog of a line. With one edge and two vertices, there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that difference. So they cannot be multiples of 3 simultaneously. In symbols: $a - b = 1$. If all vertices were multiples of 3, then $a - b$ would also be a multiple of 3. Since $a - b = 1$, it is also true that

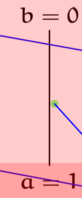

$$a - b \equiv 1 \pmod 3,$$

where the mathematical notation $x \equiv y \pmod 3$ means that x and y have the same remainder (the same modulus) when dividing by 3. In this one-dimensional version of the game, the quantity $a - b$ is an *invariant:* It is unchanged after the only move of increasing each vertex on an edge.

Perhaps a similar invariant exists in the two-dimensional version of the game. Here is the square with variables to track the number at each vertex. The one-dimensional invariant $a - b$ is sometimes an invariant for the square. If my move uses the bottom edge, then a and b increase by 1, so $a - b$ does not change. If my move uses the top edge, then a and b are individually unchanged so $a - b$ is again unchanged. However, if my move uses the left or right edge, then either a or b changes without a compensating change in the other variable. The difference $d - c$ has a similar behavior in that it is changed by some of the moves. Fortunately, even when $a - b$ and $d - c$ change, they change in the same way. A move using the left edge increments $a - b$ and $d - c$; a move using the right edge decrements $a - b$ and $d - c$. So $(a - b) - (d - c)$ is invariant! Therefore for the square,

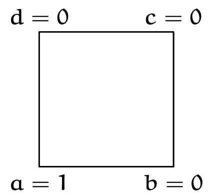

$$a - b + c - d \equiv 1 \pmod 3.$$

Therefore, it is impossible to get all vertices to be multiples of 3 simultaneously.

Comments on page 2

Are we really lowering standards? It seems like we still want correctness we are just using fewer constraints. I might prefer the term decreasing complexity? simplifying? Just the "lowering standards" has a bad connotation and there is nothing wrong with reducing the problem.

I can see how a square is helpful- isn't this simplification a little too simple to be of any help?

I get the point,but I feel like this is oversimplification. Knowing that it doesn't work in one dimension is pretty obvious, but I wouldn't consider it evidence suggesting that it does not work in 3 dimensions.

It's not supposed to show that it doesn't; it's just supposed to show how simplifications can form a relationship.

I too think that this is almost getting "too" simplistic. I can definitely understand why this example is impossible, but can't immediately see how the cube is impossible.

this is a very confusing way to say a very simple thing.

I think its ok–he says in words "they cannot be multiples of 3 simultaneously", then shows it in symbols. I understand it fine.

I think he's trying to describe with words something that is best visualized (I was able to follow along if I pictured the numbers at the vertices shifting)

I think it is kind of intuitive that they cannot both be multiples of three. If you add the same number to two numbers one apart from each other the difference between them will always be one.

I like this explanation... it does seem almost intuitive, but a bit tougher to explain. I think this explanation is very clear and convincing.

I find the explanation confusing.

But proving that the difference between a and b is always 1 doesn't apply to the whole problem since you can change a without changing b

I don't know what this is, so it'd be nice to explain it in the text.

Another form of invariance- what stays constant among variables

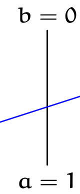

The square is easier to analyze than is the cube, but standards can be lowered farther by analyzing the one-dimensional analog of a line. With one edge and two vertices, there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that difference. So they cannot be multiples of 3 simultaneously. In symbols: $a - b = 1$. If all vertices were multiples of 3, then $a - b$ would also be a multiple of 3. Since $a - b = 1$, it is also true that

$$a - b \equiv 1 \pmod{3},$$

where the mathematical notation $x \equiv y \pmod{3}$ means that x and y have the same remainder (the same modulus) when dividing by 3. In this one-dimensional version of the game, the quantity $a - b$ is an *invariant:* It is unchanged after the only move of increasing each vertex on an edge.

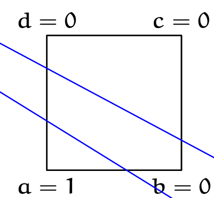

Perhaps a similar invariant exists in the two-dimensional version of the game. Here is the square with variables to track the number at each vertex. The one-dimensional invariant $a - b$ is sometimes an invariant for the square. If my move uses the bottom edge, then a and b increase by 1, so $a - b$ does not change. If my move uses the top edge, then a and b are individually unchanged so $a - b$ is again unchanged. However, if my move uses the left or right edge, then either a or b changes without a compensating change in the other variable. The difference $d - c$ has a similar behavior in that it is changed by some of the moves. Fortunately, even when $a - b$ and $d - c$ change, they change in the same way. A move using the left edge increments $a - b$ and $d - c$; a move using the right edge decrements $a - b$ and $d - c$. So $(a - b) - (d - c)$ is invariant! Therefore for the square,

$$a - b + c - d \equiv 1 \pmod{3}.$$

Therefore, it is impossible to get all vertices to be multiples of 3 simultaneously.

Could you explain modular arithmetic in class? I have a vague idea what it does, but it might be useful to go over it in class.

I think what mod tells you is the reminder, such as 5 mod 3 is 2, and 6 mod 3 is 0

True, but I haven't seen mod's in a very long time

there's nothing else to it...

I agree, I don't see why you have to write this as 1 mod 3...how does this guarantee a multiple of 3...

This statement is saying that a-b will always equal 1 mod 3, which means that a-b will never be divisible by 3. The only way that a-b could be divisible by 3, and consequently, the only way that a and b could be divisible by 3, is if a-b equaled 0 mod 3.

The reason the modulus is used is because in this problem, we only care if the numbers are a multiple of 3- it doesn't matter exactly what multiple they are, as long as they are a multiple. In other words, we don't care if the number is 3, 6, 9, etc., as long as that number always equals 0 mod 3.

Modulus just means remainder; this is the same standard term used in programming languages like Python (which is what this is all made in).

The way this is phrased is a little confusing to me, even after I read through it a few times.

Is this really necessary? I find it kind of confusing and not really helpful to the explanation of the problem. I think you should either expand on it more or scrap it entirely.

agreed, it's slightly confusing. maybe writing: since a-b=1 and 1 is the same thing as 1(mod)3, it is also true that...

This gets confusing for me because I have no idea what mod means. I get the idea from the fact that a-b will always be 1 and there's no way that you can have multiples of 3 that equal 1 when one is subtracted from the other.

This is a nice concise explanation of mod, it conveys the idea while not going into too much unnecessary detail.

As someone that has never been exposed to "modulus" I found this part of the example very confusing... Why are we using it?

The square is easier to analyze than is the cube, but standards can be lowered farther by analyzing the one-dimensional analog of a line. With one edge and two vertices, there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that difference. So they cannot be multiples of 3 simultaneously. In symbols: $a - b = 1$. If all vertices were multiples of 3, then $a - b$ would also be a multiple of 3. Since $a - b = 1$, it is also true that

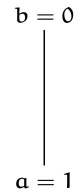

$$a - b \equiv 1 \pmod 3,$$

where the mathematical notation $x \equiv y \pmod 3$ means that x and y have the same remainder (the same modulus) when dividing by 3. In this one-dimensional version of the game, the quantity $a - b$ is an *invariant:* It is unchanged after the only move of increasing each vertex on an edge.

Perhaps a similar invariant exists in the two-dimensional version of the game. Here is the square with variables to track the number at each vertex. The one-dimensional invariant $a - b$ is sometimes an invariant for the square. If my move uses the bottom edge, then a and b increase by 1, so $a - b$ does not change. If my move uses the top edge, then a and b are individually unchanged so $a - b$ is again unchanged. However, if my move uses the left or right edge, then either a or b changes without a compensating change in the other variable. The difference $d - c$ has a similar behavior in that it is changed by some of the moves. Fortunately, even when $a - b$ and $d - c$ change, they change in the same way. A move using the left edge increments $a - b$ and $d - c$; a move using the right edge decrements $a - b$ and $d - c$. So $(a - b) - (d - c)$ is invariant! Therefore for the square,

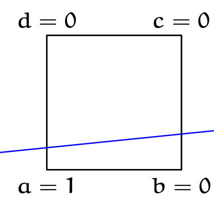

$$a - b + c - d \equiv 1 \pmod 3.$$

Therefore, it is impossible to get all vertices to be multiples of 3 simultaneously.

I recall one of the previous sections mentioned invariants without this helpful following description - perhaps this section should preceded any other mention of invariants?

So for problems like these, we should look for properties of a problem that DO NOT change after every step. IS that what invariant means?

i think that's even a stretch for this problem

Yes an invariant is some property of the problem that doesn't change after every step.

so this is where the symmetry analogy comes in.

so this is where the symmetry analogy comes in.

I actually thought the previous memo was a good introduction to invariants, and then this memo followed up quite nicely with a longer example.

really like the progression from line to square to cube. Helped me understand the problem a lot better.

This is just a nick-picky thing (or just a lazy thing), but this paragraph is not hard to follow; however, I didn't really read it and refer to the diagram - I just assumed I knew which edges were being chosen...another small graphic, or even color could help with this.

meh you don't really need to know.

I know what you mean... I didn't feel like slowing down and mentally going through the proof. However, when I got to the end of the 3 pages, I wasn't sure why the argument was proved so I decided to actually examine the proof. And, it did make sense. But I feel like there is a psychological aversion to proofs, especially more mathematically daunting ones which make people avoid reading paragraphs like these. The one difference between the examples in this reading versus the ones in the previous reading is that the previous reading had it organized and spaced out better. The paragraphs here are long and have equations/symbols mixed right in. It would be much easier to read a proof, if the equations were spaced out on separate lines with less text in between new lines.

The square is easier to analyze than is the cube, but standards can be lowered farther by analyzing the one-dimensional analog of a line. With one edge and two vertices, there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that difference. So they cannot be multiples of 3 simultaneously. In symbols: $a - b = 1$. If all vertices were multiples of 3, then $a - b$ would also be a multiple of 3. Since $a - b = 1$, it is also true that

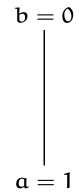

$$a - b \equiv 1 \pmod 3,$$

where the mathematical notation $x \equiv y \pmod 3$ means that x and y have the same remainder (the same modulus) when dividing by 3. In this one-dimensional version of the game, the quantity $a - b$ is an *invariant:* It is unchanged after the only move of increasing each vertex on an edge.

Perhaps a similar invariant exists in the two-dimensional version of the game. Here is the square with variables to track the number at each vertex. The one-dimensional invariant $a - b$ is sometimes an invariant for the square. If my move uses the bottom edge, then a and b increase by 1, so $a - b$ does not change. If my move uses the top edge, then a and b are individually unchanged so $a - b$ is again unchanged. However, if my move uses the left or right edge, then either a or b changes without a compensating change in the other variable. The difference $d - c$ has a similar behavior in that it is changed by some of the moves. Fortunately, even when $a - b$ and $d - c$ change, they change in the same way. A move using the left edge increments $a - b$ and $d - c$; a move using the right edge decrements $a - b$ and $d - c$. So $(a - b) - (d - c)$ is invariant! Therefore for the square,

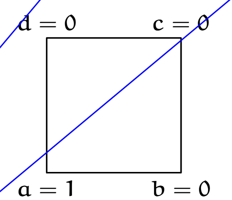

$$a - b + c - d \equiv 1 \pmod 3.$$

Therefore, it is impossible to get all vertices to be multiples of 3 simultaneously.

I definitely get what this is saying, and can picture how the different combinations may or may not change a-b or d-c, and I like how it connects with the invariant we were discussing with the one dimensional line. I just kind of have this strange uneasiness with some of the phrasing. I can't put my finger on it, but since I understand, I guess it might not be such a big problem.

I was struck by this same feeling. I read the sequence a few times trying to pinpoint it with no luck. I dont know if diagrams or equations would help or if its just something that takes a few reads to swallow

Wow. That was pretty eye opening. Can we get a little clearer of an explanation in class? I think I understand it, but I also think I am confused enough to not be able to explain it to someone else.

Woah, that's really neat, but I don't think I would have ever figured that out on my own! I guess before I read ahead I'll try to apply it to the cube...

I think a good way to think about this is that every edge touches either a or c, and every edge touches either b or d (but never both). So each edge adds 1 to the value of (a+c) and to (b+d). Hence (a+c)-(b+d) is invariant.

it'd be helpful to see all 3 invariants written out separately on the left or something so we can see more clearly how they come together to get the concluding eqn.

Actually the explanation you just gave is really a great one. I think its much clearer than the one given in the text.

This is extremely helpful. Taking advantage of the fact that every edge touches a or c or b or d but not both makes the problem solvable.

Very neat. I was wondering how the overly drawn out explanation of the line earlier would apply. Very cool.

This is where solving the problem becomes clear. I was a bit shaky before this.

I agree. wow. this is pretty crazy. that (a-b)-(d-c) stays constant. this makes solving the problem so much clearer

I don't really see how this problem demonstrates symmetry. It seems more like the point is invariants, and while often invariants and symmetry are linked, I don't think that link is very clear in this example.

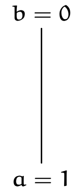

The square is easier to analyze than is the cube, but standards can be lowered farther by analyzing the one-dimensional analog of a line. With one edge and two vertices, there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that difference. So they cannot be multiples of 3 simultaneously. In symbols: $a - b = 1$. If all vertices were multiples of 3, then $a - b$ would also be a multiple of 3. Since $a - b = 1$, it is also true that

$$a - b \equiv 1 \pmod 3,$$

where the mathematical notation $x \equiv y \pmod 3$ means that x and y have the same remainder (the same modulus) when dividing by 3. In this one-dimensional version of the game, the quantity $a - b$ is an *invariant:* It is unchanged after the only move of increasing each vertex on an edge.

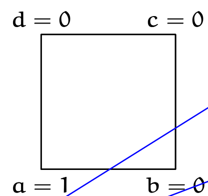

Perhaps a similar invariant exists in the two-dimensional version of the game. Here is the square with variables to track the number at each vertex. The one-dimensional invariant $a - b$ is sometimes an invariant for the square. If my move uses the bottom edge, then a and b increase by 1, so $a - b$ does not change. If my move uses the top edge, then a and b are individually unchanged so $a - b$ is again unchanged. However, if my move uses the left or right edge, then either a or b changes without a compensating change in the other variable. The difference $d - c$ has a similar behavior in that it is changed by some of the moves. Fortunately, even when $a - b$ and $d - c$ change, they change in the same way. A move using the left edge increments $a - b$ and $d - c$; a move using the right edge decrements $a - b$ and $d - c$. So $(a - b) - (d - c)$ is invariant! Therefore for the square,

$$a - b + c - d \equiv 1 \pmod 3.$$

Therefore, it is impossible to get all vertices to be multiples of 3 simultaneously.

I'm confused with 1(mod3) notation. Is the "1" showing that the remainder has a difference of 1?

I'm also a little confused with this, maybe it's because I don't entirely understand mod()

This means that the answer to (a-b)+(c-d) =1 when modded by 3. For example, this means that the only possible answers to (a-b)+(c-d) are 1, 4, 7...

Thanks for the explanation I would have had the same question.

But a-b+c-d = 1, exactly and always, not just 1 (mod 3)...

We only care about it being 1 (mod 3) because our goal is 0 (mod 3) != 1 (mod 3).

Would an explanation involving linear combinations help? Sorry, after taking 18.06 this math seems very applicable to spaces and linear combinations.

Even though we've found this, it was obviously much more difficult to find than the 1-d case- was there no way of generalizing to make this easier for a cube?

This was much easier for me to solve by taking (a+b+c+d) = 1 mod 2. Any move you make increases two vertices by 1, so the above statement is always true. Unfortunately, the solution has the sum equal to 0 mod 2.

However, using mod3 is helpful in this case because we are trying to make all edges equal to a multiple of 3... I see how the way you did it is equivalent, but for someone without much knowledge of mods it is useful to see the original problem statement in the equation

we don't know that a solution has to have a+b+c+d = 0 (mod 2). What if a possible solution were a=b=c=3, d=6? That sum is 1 (mod 2), and it's not obvious from a (mod 2) argument that it's an impossible configuration

That wasn't as enlightening/clarifying for me as it seems to have been for other people.

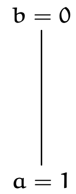

The square is easier to analyze than is the cube, but standards can be lowered farther by analyzing the one-dimensional analog of a line. With one edge and two vertices, there is only one move: incrementing the top and bottom vertices. The vertices start with a difference of one, and continue with that difference. So they cannot be multiples of 3 simultaneously. In symbols: $a - b = 1$. If all vertices were multiples of 3, then $a - b$ would also be a multiple of 3. Since $a - b = 1$, it is also true that

$$a - b \equiv 1 \pmod 3,$$

where the mathematical notation $x \equiv y \pmod 3$ means that x and y have the same remainder (the same modulus) when dividing by 3. In this one-dimensional version of the game, the quantity $a - b$ is an *invariant:* It is unchanged after the only move of increasing each vertex on an edge.

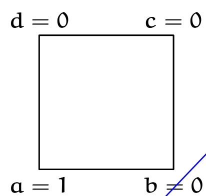

Perhaps a similar invariant exists in the two-dimensional version of the game. Here is the square with variables to track the number at each vertex. The one-dimensional invariant $a - b$ is sometimes an invariant for the square. If my move uses the bottom edge, then a and b increase by 1, so $a - b$ does not change. If my move uses the top edge, then a and b are individually unchanged so $a - b$ is again unchanged. However, if my move uses the left or right edge, then either a or b changes without a compensating change in the other variable. The difference $d - c$ has a similar behavior in that it is changed by some of the moves. Fortunately, even when $a - b$ and $d - c$ change, they change in the same way. A move using the left edge increments $a - b$ and $d - c$; a move using the right edge decrements $a - b$ and $d - c$. So $(a - b) - (d - c)$ is invariant! Therefore for the square,

$$a - b + c - d \equiv 1 \pmod 3.$$

Therefore, it is impossible to get all vertices to be multiples of 3 simultaneously.

I don't quite follow the implications of this therefore. Why does this expression prove that solving the problem for the square is impossible?

This took me a while, and this isn't the most rigorous answer, but if all the numbers were multiples of 3, then a-b+c-d would also be a multiple of 3 (could be positive or negative or 0). The fact that it's offset by this remainder of 1 means that at least one of the numbers isn't a multiple of 3.

Ahhh thank you for that short explanation, I wasn't sure how we got to the therefore either

Yea, he's stating the invariant, which is that a-b+c-d is 1(mod 3). Invariant means this statement will always hold no matter what move you make in the game, so that's why after stating it once he can say therefore, etc.

So the earlier guess was right, and this makes sense

Yeah, I am definitely surprised by this answer since it didn't agree with my initial gut, but this makes sense now why this is true.

I thought of this another way originally. The sum of all vertices increases by 2 each time, and for all 8 to equal 3 the total sum must be an even number (24). But 1+ even numbers is always odd, so at no point can't it sum to 24.

Is it possible for other numbers?

but does this hold true for the cube? the cube is a fundamentally different problem

my way seems a lot simpler than this. Is it correct?

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

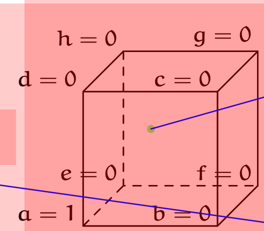

$$a - b + c - d + f - g + h - e \equiv 1 \pmod{3}$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

Comments on page 3

Are there confusing varieties that are solvable with only one solution?

I think this is really cool, but I still don't think I could ID and solve one of these on my own (the physics problem below excluded, because that one made sense beforehand)

So this is where I got stuck really fast – I wasn't sure how to generalize it, and what still held true, since I was so confused with all the vertices. If there a good way to essentially put together the problem again once it's been broken down?

Try to find the sets of vertices that every edge contains exactly 1 of, for example <a, c, f, h>. The other set is <b, d, e, g>. This is basically how the previous cases were formed, although it wasn't stated explicitly. Now, each edge you add will add 1 to exactly one element of each set... and the invariance can be proved from there.

Adding this explicit explanation to the text would be great. I subconsciously did this, but hadn't figured out why. You can see the pattern in how the letters alternate between +/-, but this could depend on how you assign the letters to the vertices.

yeah, having a more detailed explanation about how we got to this equation would be very helpful...while I understood how to get the equations for the line and the cube, I kind of got confused in this last part, which is the most important to solve the problem!

This is really cool, but what is your general approach for finding invariants? In general, I think I would have a tough time computing what the invariant in a problem would be.

how did you choose which direction the vertices were chosen? Or does it not matter, so long as they lie on two different squares?

See below for a good explanation (tied to the box that starts around "this impossibility").

Yeah, how were the edges chosen? The front face has horizontals chosen and the back face has verticals chosen. Is this crucial?

I would like to be walked through this in class.

I think I would actually have tried this first and probably subconsciously done the divide and conquer, but I really like how it's being fleshed out here.

How did you prove this invariant is true for the cube?

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

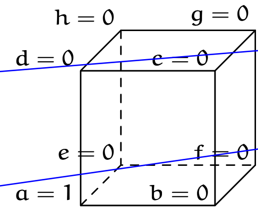

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

i find the problem to be spatially ambiguous in some respects and have a hard time visualizing all the moves.

I am still a little confused on this equation. Can you go over it in lecture?

The previous example where we "lowered our standards" really helped me to understand this more complex problem. In this way, I can relate this problem to my understanding of symmetry as well as divide and conquer.

I was right (referring to the question posed in the beginning)!

haha, me too! For once I wasn't outdone by some clever shortcut!

I'm kind of disappointed. I was hoping for a really cool shortcut to solve this problem.

Same here, because my intuition told me there wasn't a solution I was ready to be proved wrong by a simplification I never would have seen.

I never would have thought to look at this problem this way, but it is incredibly useful now that I see it.

Me too–I've never actually thought of this as a problem solving strategy. see what stays constant, so you know what can change relative to what. applies to much more than this problem.

Just an FYI... I believe both 6.042 and 6.005 use the term invariant and I remember there being some confusion between the two as they used slightly different definitions. So it might be helpful to formally define the term earlier as I believe some students have never heard the term before, and other students have heard too many definitions.

I agree that a formal definition would be helpful. "quantities that remain unchanged" seems like it might be misleading when your invariant is a relationship between systems that can't easily be explained with a numerical constant.

I think this concept should be put at the beginning of the section so the reader kind of knows what to look for when thinking about the problem.

I agree. It would be very helpful in understanding what you meant in the first page when you mention invariants if we were given a definition earlier in the section.

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

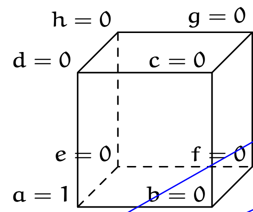

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

I like this problem because it explains the notion of invariants in a different context than explained in lecture. Now I see their usefulness.

great quote

Forgive my ignorance, but I'm not sure how these are "solitaire" games.

Yes, how exactly are you defining solitaire? Isn't any problem solved by oneself a solitaire game?

...unless you collaborate

lolz

I think "solitaire games" refer to a subset of games in which the result of the game depends only on your actions and not on another's.

solitaire games are all games with some invariant?

solitaire in that only one person plays, right, not like this models a card game?

puzzles! not games...

I assume you mean an unpowered roller coaster *car*?

or maglev!

I'm pretty sure I've seen this before, and spent a good long time trying to find a way to lift the roller coaster above its starting point! The only thing I think I could come to was to violate conservation of mass...and dump people out

yeah right away this sounds fishy.. im pretty sure physics doesnt work this way

Is this even possible? (doesn't conservation of energy say it's not, even with lots of loops and curves?)

Exactly – that's why energy is a convenient invariant in physics.

I still think the idea of throwing people off the coaster after you've gained speed is a pretty solid one.

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

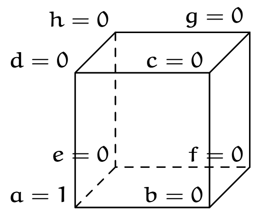

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

The rollercoaster example really helps make the point. it's something that we all can understand, so it shows how applicable this principle of looking for invariants really is.

Using GIR-level physics is great for explaining these concepts, especially right after an example that may have lost some of the less engineering-inclined.

this sounds awkward

agreed

Overall clear, easy to understand, and makes a good point. No real comments here.

This is a fresh look at a concept i've known since high school. Looking at different properties in physics or engineering can simplify a lot of problem... or at least make them more understandable.

this is also found in LTI systems, and optical systems

The inclusion of invariants that we've all seen before seems to be an effective way to tie it all in together. Its helped a lot whenever that has happened in the book so far

This would have clarified how to solve the cube problem if I had seen this first

I agree, the first time an invariant was introduced it would have been nice to see a type of invariant like this that everyone knows.

I can kind of understand the connection between this example and what we are learning, but not entirely. Could you explain more in lecture?

yeah that makes more sense

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

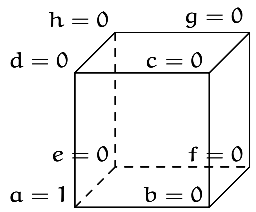

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

I feel like this example is a little too obvious. I'd be surprised if there was anyone in the class that didn't know this right off the bat. I feel like a more difficult real-life example would be more appropriate here.

I don't know - while the answer was obvious, it is kind of nice to have a problem that I know the answer for sure. It is actually great, because it really gets you to think about and understand the concept of the invariant point, because it points out that you've been doing it all along!

I agree. All the examples up to this point have been kind of non-trivial. We know how to solve this problem, but that allows us to pay attention to the invariant part as opposed to solving the problem itself. This whole section in general I feel really helped solidify what was meant by "When there is change, look for what does no change"

I feel like energy is also a sort of abstraction, it hides all of the more complicated underneath it and allows us to greatly simplify most real world problems.

Agreed - it's nice to see something from old 8.01 for some familiarity.

I think this is a great example because it's simple, clear, and easy to understand. An additional problem that is more difficult would be fine too, but as a supplement, not a substitute.

My comment would be to reorder examples. It's really hard to follow math in text and I think the first section needs to be hashed out more. Seeing this example first as a "warm-up" might make me more prepared to understand the first bit.

I would see what it is like reordered; I also think that the energy example is easier, but I just spent some time looking at the cube example. It may be better reordered to get the brain around the concept early.

Actually, I like the ordering – I struggled a lot with the first one and then came back to this and it helped really solidify the concept seeing something familiar. If it'd been ordered hte other way around I don't think I would have gotten that same "It makes sense!" moment, since I'd be too lost in the details...

I agree with that. To me, I only really get the topic in this class if I couldn't solve it using my old methods. I struggle to get everything out of this example because I want to revert back to what I already know.

I think that putting the physics example first would have helped to hit home the point that one way of solving difficult problems is to find invariants.

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

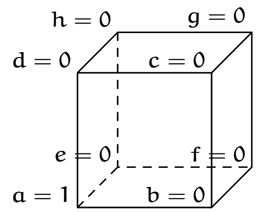

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

yeah I agree that having the more simple example first would be helpful in understanding the invariant method...also I disagree that this example is too obvious, while it is simple to understand, I haven't seen 8.01 stuff since high school so it wasn't that obvious!

I agree that this example would've been a nicer introduction to the invariant, and the cube being more an application than an introduction.

Seeing something familiar gives me faith in the approach. After finishing the reading, the last problem forced me to go back and make more sense of the cube problem because I knew there was truth in the approach.

I think it would be nice to have an example that was this simple and short toward the beginning of the section to introduce the idea of an invariant and make the section easier to follow.

While I am unsure as to whether or not the examples should be rearranged, I really feel that this example solidifies this section.

It's nice to see something I know but I feel like I've been told energy is an invariant since high school. This doesn't provide practice with finding invariants; although it does show it's relevance. I would like another "harder" example in this chapter.

This is exactly what I thought of doing after I read the problem. Sounds like I'm getting the hang of this!

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a-b+c-d+f-g+h-e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

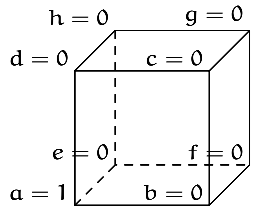

$$a-b+c-d+f-g+h-e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

Now what if the car split in half? Then you could get each half of it above the starting height (at different times) without violating the conservation of energy. I believe you would add entropy though...

Really? Would you be able to do that?

split it in half or raise each half above the starting height? If splitting the train in half took no energy (perhaps you could just uncouple a car), then it would not violate the law of conservation of energy to have one rise and the other fall. Suppose two cars were strung together with a cable, and the cable was looped over a pulley. One car could fall and the other rise, and since friction is negligible, the process would be reversible.

I feel like the issue with this is that you have to put energy into the system to separate the cars. Without friction, if you split both the cars in half both halves will continue to travel in the same direction at the same speed together. The only way to separate them is to use some device that pushes them apart, and that would require adding energy to the system.

I agree. Moving this forward would really help to hit home the point of the section. Which is that looking for invariant quantities can simplify problem solving tremendously.

Not so. The energy required to split them needn't be wasted if it comes from something conservative like a spring. As long as at the end of the day both cars have v=0 and the height of their center of mass is at hstart, you can rearrange their parts any way you want without expending any net energy (assuming no friction etc.)

I understand how identifying invariants is critical to making a symmetry argument, but I'm still a little confused how the cube is an example of a symmetry problem? Is it simply because we were able to generalize from the 1D/2D cases? How exactly is this different from abstraction or divide and conquer?

The symmetry doesn't come from the reduction of dimensions, but rather from the fact that no matter which move you make, there is some quantity (the invariant) which is preserved. This is just like saying that no matter how many 90 degree rotations you apply to a square, its shape is the same. The reduction to 1D just helped us discover the invariant behind the symmetry.

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

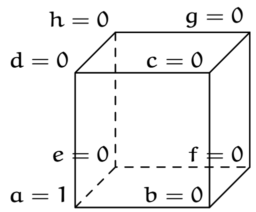

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

Maybe this is a good time to tie in Noether's theorem. There's probably not enough space to prove it, but it sure shows how powerful symmetries are in physics. (Noether's theorem says that for each symmetry there is a corresponding conserved quantity. Time-, translational-, and rotational-invariance correspond to energy, linear momentum, and angular momentum conservation laws.)

umm i dont think thats necessary. we got the point.

I had never heard of Noether's, and it sounds applicable, but maybe at too high a level for right here. It could be recommended for deeper scope.

I think stating this, or something similar to this, in the introduction would help clarify the purpose of this section. As it is it just seems like a long example with it's relevance not known until the end.

Agreed. I found that reading the section a second time was easier, since I constantly referred to this statement.

similar idea to recursion–looking for patterns inside patterns, ie patterns that stay constant throughout the large pattern

great summary.

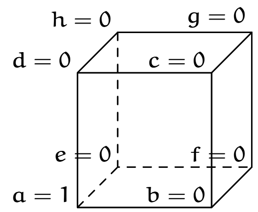

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

$$a - b + c - d + f - g + h - e \equiv 1 \pmod{3}$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

I like paragraphs like this, it summarizes everything so concisely yet easy to understand

Agreed

Yeah they're really helpful.

Is it intentional irony that the moral didn't change?

It's appropriate, not ironic.

I also liked how the ending paragraph ties together multiple units (symmetry and abstraction)

So truly, each subsequent section does build on each other.

I agree, I was getting confused as to how this was going to help us make estimations. this section was more about solving this logic puzzle.

I concur; in fact, repeating exactly what was taught in the previous section with another example that specifically stresses the point via invariants has hit the nail on the head so to speak.

I agree with all... I was sorta wondering how symmetry would be brought up... or simply how this applies to what we've been looking at. Everything is brought together nicely in the paragraph though. Very cool.

This paragraph is great; The bringing together of our estimations tools as building blocks really helps me to see how things can be applied. These things are not just tools we can do down a check-list for, they can be combined!

On a side note to symmetry, there are also cool theorems in group theory that say for example that there are only 17 different ways to make wallpaper, and only 7 ways to make a strip of horizontal wallpaper, all of this proved by symmetry arguments. But definitely beyond the scope of this section, maybe a cool one liner somewhere.

I think the last paragraph of the readings is always a very helpful summary.

I'm confused- I thought this chapter was about symmetry, why is the conclusion talking about only abstraction. To be honest I'm not clear on what is meant by symmetry.

The original three-dimensional solitaire game is also unlikely to be winnable. The correct invariant shows this impossibility. The quantity $a - b + c - d + f - g + h - e$ generalizes the invariant for the square, and it is preserved by all 12 moves. So

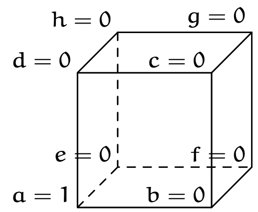

$$a - b + c - d + f - g + h - e \equiv 1 \pmod 3$$

forever. Therefore, all vertices cannot be made multiples of 3 simultaneously.

Invariants – quantities that remain unchanged – are a powerful tool for solving problems. Physics problems are also solitaire games, and invariants (conserved quantities) are essential in physics. Here is an example: In a frictionless world, design a roller-coaster track so that an unpowered roller coaster, starting from rest, rises above its starting height. Perhaps a clever combination of loops and curves could make it happen.

The rules of the physics game are that the roller coaster's position is determined by Newton's second law of motion $F = ma$, where the forces on the roller coaster are its weight and the contact force from the track. In choosing the shape of the track, you affect the contact force on the roller coaster, and thereby its acceleration, velocity, and position. There are an infinity of possible tracks, and we do not want to analyze each one to find the forces and acceleration.

An invariant – energy – vastly simplifies the analysis. No matter what tricks the track does, the kinetic plus potential energy

$$\frac{1}{2}mv^2 + mgh$$

is constant. The roller coaster starts with $v = 0$ and height h_{start}; it can never rise above that height without violating the constancy of the energy. The invariant – the conserved quantity – solves the problem in one step, avoiding an endless analysis of an infinity of possible paths.

The moral of this section is the same as the moral of the previous section: *When there is change, look for what does not change.* That unchanging quantity is a new abstraction (Chapter 2). Finding invariants is a way to develop powerful abstractions.

I think this reading as a whole was the best so far. Interesting and easy to understand throughout. There wasn't a point where I got stuck because of wording or complexity, and there were some great quotes to keep me reading.

I also really enjoyed this reading. The presentation of a complex, non obvious example followed by an obvious example really highlighted the points well.

I really dont like puzzles...this was a good way to look at solving this one...but i still don't like puzzles.

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

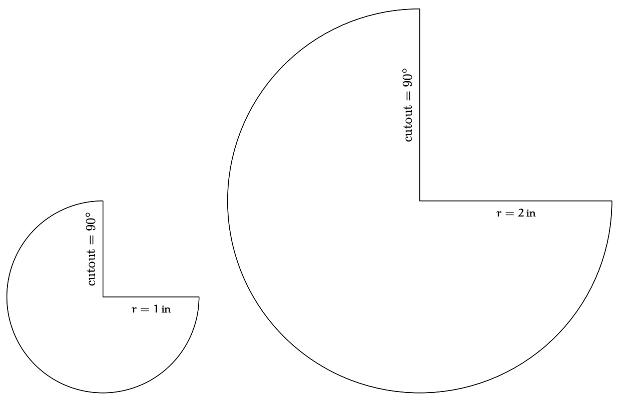

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

But the cross-sectional area of the cone changes...do you just use the median diameter or something to measure A?

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

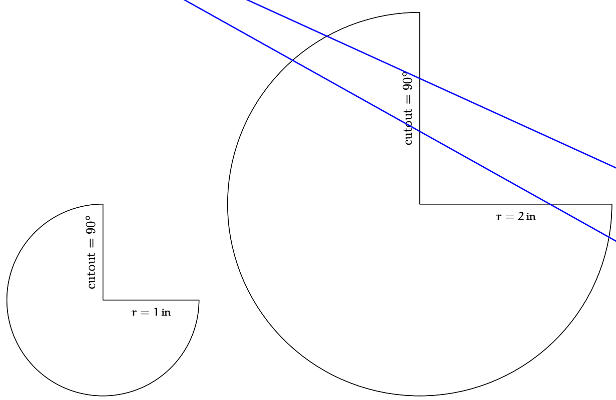

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

Read Sections 3.3 (drag) and 3.4 (application to cycline) for the memo due Tues at 10pm. See you Wednesday (no lecture on Monday).

why is drag one of the most difficult subjects in classical physics?

Nice to get the explanation of this.

from the first day of class!

Yeah, this is definitely a nice throwback to the early days of class lecture.

instead of using a sub-heading, I'd use an inset box around the experiment instructions.

How much of a difference does this make? Are they just too unstable at smaller sizes?

and they are easier to cut out if they are larger

Generally easier to tape and handle. Any imperfections in cutting them out technically have a larger impact because of the percentage in size with respect to the entire shape, but this is likely minimal.

This makes it extremely unlikely that anyone is going to do this (especially if this becomes an actual textbook). It might be better to have something that can be traced (in an appendix if it takes up too much room) or just give the specifications and assume people can make a decent attempt at drawing a circle.

typo, only one of these words is needed

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

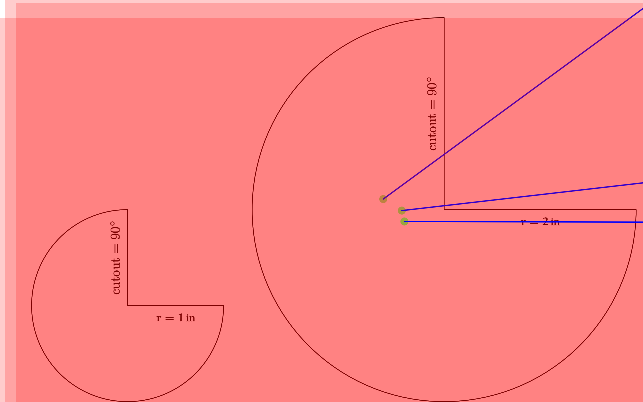

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

Do the words on these templates matter? I had to zoom in pretty far to see what they said.

I think he just wanted to make clear what the radii of the two were, but I agree. He should have written that somewhere in the text instead of on the tiny picture.

or just enlarge the text?

well when you copy it it's supposed to be at 200% enlargement.

regardless, i don't know if the words would come out clearly if you were to magnify the shapes.

it might become pixelated when you enlarge it. i think it's so small because he used a much larger picture to start with and latex'd it in without realizing that the original text would get shrunk so small.

I would think that the smaller one falls more slowly. It seems like this is the case for animals.

Yea we shouldn't have to actually do the experiment because we watched it in class the 1st day. Ratio 1:1 about

Don't forget, this is a book intended to for public use, so not everyone will have seen the experiment. Though it would be nice if the images were given a separate section so I wouldn't have to worry about enlarging.

This was the experiment we did in the first lecture right?

That's right: I am delivering one of the promises made so far.

Hooray for promise fulfillment!

Haha yeah, if I recall, the ratio was roughly 1:1 right?

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

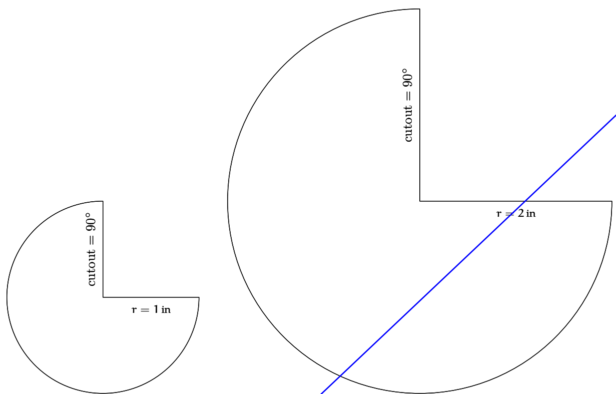

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

There's never any discussion of how the cone's effective area differs from the area of the circle or the circle minus the cutout. (and what if the cutout had not been 90 degrees?)

http://math.about.com/od/formulas/ss/surfaceareavol_2.htm

Basically, area = pi*r*s+pi*r^2

True, but I think at this first pass, it's sufficient just to notice that it's larger so that we can set up the competition between increasing force and increasing drag.

Yeah–there is no mention about how the cone's effective area relates to the area of the cutout itself, but we don't need the exact relations, we just need the ratios. We know that drag is proportional to area, and we assume both cones are cut and made the same way (only difference being size). Thus, we know the ratios of radii, which we can then do math on to find the ratios of the effective areas

It might be useful to define what you consider drag to be before using it in a sentence.

Agreed. I remember that drag is the force that slows things down while falling or moving fast, but nothing more than that. Unless you are intentionally keeping it vague for now.

It might also be helpful since in the next paragraph the text refers to the "drag force" instead of simply "drag". I think an explicit definition would make this transition a bit more clear.

I feel like people have an intuitive grasp of what drag is - I've never taken a fluids class and the concept of drag being some force that makes you go slower isnt that difficult to grasp. I think it comes down to a question of intended audience and what they should know

Agree to define what drag is, people who are not course 2 or course 16 might have just heard of it but never have worked with it in formulas.

To me, this reads very awkwardly.

It would be useful if you explained how these two things play into the terminal velocity of each of the cones.

I agree with the awkward wording sentiment. Instead of "win" maybe is larger?

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

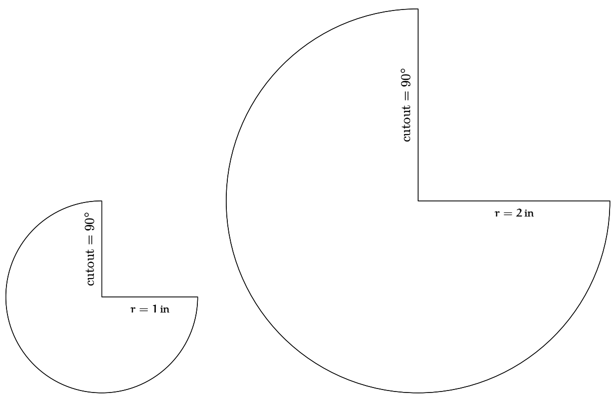

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

at this point I would have stopped and done the experiment (if we hadn't already done it in class).

i think that the section would be a lot better if you referenced the outcome after this...most people aren't going to wait until the end of the section to try it out (at least the ones I know).

ie. Through this experimentation you find that the two cones fall at roughly the same speed. However, actually calculating the drag forces to prove this is a very complicated. It requires solving...

But no one would actually use this formula. We would estimate based on the surface area!

So given this is a continuous book, the N-S equations have been shown 2-3 times by now, and I think it's unnecessary to show it this many times.

I disagree, I think you need to show the equation to illustrate how complicated it is and why approximations are useful. I think it's good to remind the reader what it looks like every time it is mentioned, even if you don't go into great detail about it every time.

Also, is it really necessary to go directly to navier stokes to solve for drag? Can't you use newton's second law and the coefficients of drag that have been approximated for various shapes to solve for the drag force?

I think the answer is yes, you can approximate this using drag coefficients and Newton's Laws, however the point is that this is still an approximation for the drag force, and not actually a calculation of the drag force itself.

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

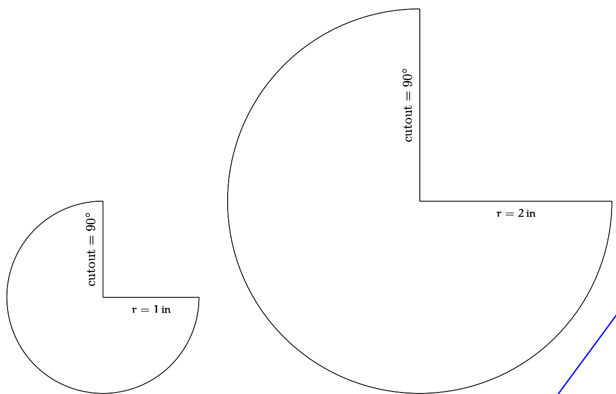

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

I really like your examples that relate to equations or principles I've used in my other Course 2 classes – it helps clarify a lot of things.

Aren't the Navier-Stokes equations unsolvable (at least at the present moment), thereby necessitating estimation or numerical methods?

Not quite. You can solve them for certain very simple cases, but I think they are only solvable in 3 or 4 situations.

Or rather, they're solvable when you can cancel out enough of the terms (which is true for certain simple cases)

Is this the same equation mentioned earlier as an example of super complicated things easier investigated via abstraction techniques?

The very same one.

What are the assumptions for the cases that you can solve N-S for again. We learned it in 2.005 but that was so long ago?

nice tie to earlier notes

Although I know the equation, it would be cool if you defined the variables so non-engineers knew what you were writing about

Yeah I agree..I've never seen this equation and I was a little confused. There is no need to explain where the equation comes from or how it derives, but perhaps just stating what each variable represents would be helpful

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

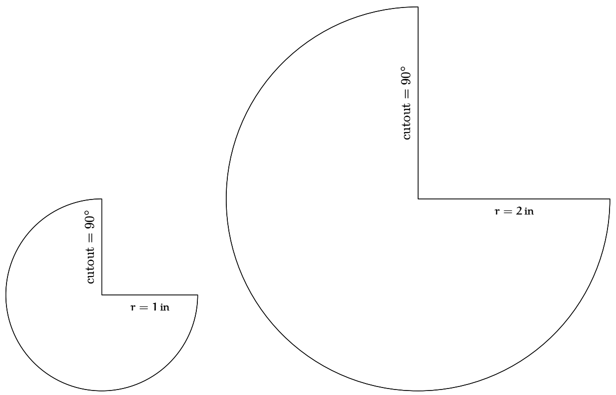

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

That sort of equation just sounds terrifying. I have rarely seen such a string of scary math words strung together like that.

Now you know why all us course 2 kids love 2.006 so much.. (I actually think its material is really cool, just really hard, too)

I'm not familiar with how to solve any of the properties of that equation... yay for abstraction and estimation!

I'm sure it's intentional ;-). And just imagine what that reads like to someone who isn't used to the sea of jargon we're already floating in at MIT.

I agree i like this sentence - makes you really want to find a way to get away from that equation!

I also think that the paragraph could have ended here. It gets kind of into the realm of irrelivnt after this.

I don't think there is a need to "drag" this out. If the point was to show that the calculation is really complex, I think even just putting only the "full description" in one go, instead of separating the equations, would do it.

3.3 Drag using conservation of energy

Conservation of energy helps analyze drag – one of the most difficult subjects in classical physics. To make drag concrete, try the following home experiment.

3.3.1 Home experiment using falling cones

Photocopy this page at 200% enlargement, cut out the templates, then tape the their edges together to make two cones:

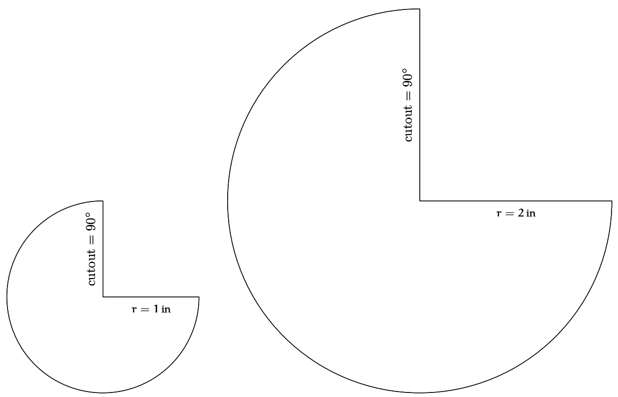

When you drop the small cone and the big cone, which one falls faster? In particular, what is the ratio of their fall speeds v_{big}/v_{small}? The large cone, having a large area, feels more drag than the small cone does. On the other hand, the large cone has a higher driving force (its weight) than the small cone has. To decide whether the extra weight or the extra drag wins requires finding how drag depends on the parameters of the situation.

However, finding the drag force is a very complicated calculation. The full calculation requires solving the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}.$$

And the difficulty does not end with this set of second-order, coupled, nonlinear partial-differential equations. The full description of the situation includes a fourth equation, the continuity equation:

What does this question mean / why is it relevant? Sorry if this is elementary but I haven't dealt with drag before. I guess it just serves to prove that solving this problem via math is quite difficult and through this mathematically analysis it is very apparent.

I think that was the point. We don't want to solve for drag using these equations, we would rather solve it using some handy estimation tool! That being said, the continuity equation is not terribly scary.

The continuity equation is useful in solving the Navier-Stokes equation as it establishes quantities that are conserved – in this case, mass (you don't lose mass!). But I agree...the point is that it is very complex (and sometimes impossible) to solve N-S.

Agreed - I think this examples is included just to demonstrate how hard the task would otherwise be.

I am familiar with these equations, but I'm not sure what the point of showing the continuity equation is...

The point of these complicated equations isn't to actually use them to solve the problems, or even for us to learn them to use in the future. The whole purpose of this class is learning ways to approximate things without using these crazy math/physics/whatever equations.

$$\nabla \cdot \mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

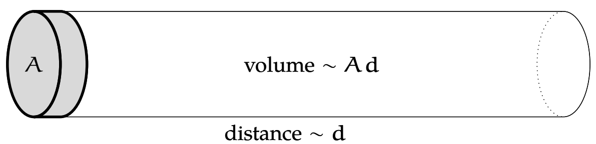

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

This pagebreak confused me - took a minute to discover that this _was_ the continuity equation

Same here...dont know if there's an easy way to correct this formatting so that the equation is close to the text that refers to it rather than hanging on the next page

I agree too. Especially if I were reading in a textbook, it would be confusing. I think when all relevant materials are in one page, it is the easiest to follow what's going on.

I agree too. Especially if I were reading in a textbook, it would be confusing. I think when all relevant materials are in one page, it is the easiest to follow what's going on.

Fluid? I thought we were dropping it in air. Can you explain why we are considering a fluid instead of considering air, or how they are related, or something along these lines?

Navier-stokes is effectively the fluid dynamics equation. It models how fluid moves and behaves in different situations (stresses, pressures, velocities, sizes of tubes), and air can be treated as a fluid (and is) when considering drag. I think a handy analogy, for this paragraph at least, would be to charge/current and maxwell's equations

It's really easy to understand air as a fluid with just a very very low density.

In fact, if you go to a high enough pressure for a given fluid, the transition between gas and liquid (both fluids) is non-existent. i.e., by changing the temperature you might change from one to the other, but there is no defined boundary (or even a difference) between the two.

I still think it would be worth it to mention that by fluid you actually mean air here.

The term "fluid" can refer to any liquid or gas that has the ability to take the shape of it's surroundings. Since the N-S equations deals with pressure and density, then we don't have to necessarily use air.

$\nabla \cdot \mathbf{v} = 0.$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

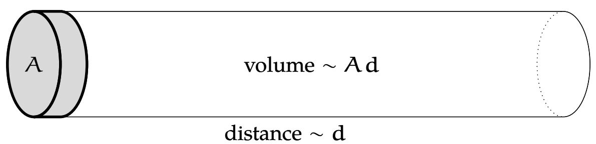

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

On first reading, I tried to go back and think through the equation, b.c. I've never had to use the navier stokes eq. But I realized that wasn't the point; the point is that it is difficult.

Yeah, that whole paragraph was kind of in one ear and out the other for me. It looks like the typical difficult textbook sentence, which is difficult to understand/absorb.

I think people can gather from looking at the equation that it is very complicated and nobody wants to ever have to solve it. Maybe this paragraph could be shorter so that it doesn't distract people from its purpose, which is to show that a quick estimation is much more enjoyable than cranking out the navier-stokes equation.

I agree that as someone who has never encountered the Navier Stokes equation before this paragraph is mostly just clutter. All I really needed was the sentence, "In short, solving the equations analytically is zero". But if you had seen Navier Stokes before it might serve as a useful review of the difficulties inherent in solving the equations.

That's an understatement.

You should–it might win you a million dollars.

why?

yeah I've heard that fluid drag is extremely complex and it is very difficult to model. Why is that? How accurate are current models?

I think those equations are only valid for "ideal" conditions, but in the real world, there are so many external factors we might not be aware of, or this equation might not be able to account for, so this makes the math very complicated and it's kinda hard to set the boundary conditions as well

$$\nabla \cdot \mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

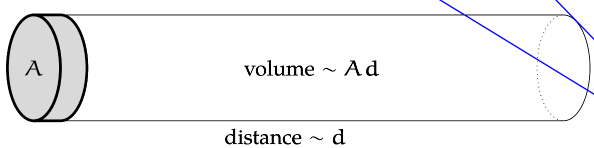

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

Does this refer to solutions for the Navier-Stokes equation? Or to solving the two equations together?

Solving them together. Navier-Stokes is just part of the system of equations that defines the physical system.

If there are only limited solutions, do MEs typically just use numerical methods to solve them?

Also, how did people design bridges before computers. For example, the Golden Gate Bridge in the SF Bay Area was designed in the 1930s, pre computers. Despite this, it is still designed to survive the high winds/etc. How did they manage to do this with these complex equations

define please! we can make a good guess of what it means using context, but it'd be nice to see this word defined specifically.

I'm also unfamiliar with this term. I'm also eager to know what 'dry water' is and how it got this name.

I'm going to guess, but water with no drag and viscosity?

yeah...inviscid means no viscosity, which helps simplify some problems when you can consider boundary conditions separate.

I think it's pretty clear what this means - coming right after discussion of zero-viscosity fluids, we should be able to infer the meaning of inviscid.

Yeah, I don't think it needs to be defined because you can guess what it means after reading the sentence.

What's this citation referencing?

interesting name

why does this matter?

why does an inviscid solution predicts 0 drag?

$$\nabla\cdot\mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

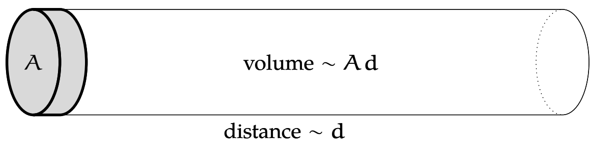

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

Can you really call it lying if it arrives a correct answer?

I'm not sure where the skillful lying comes into play.

I believe that the skillful lying is in pretending that the object is shaped a certain way in order to get the physics to work the way that you want them to.

I like this term; I think it has a nice ring to it, and adds comedy to truth. Another memorable tidbit to keep the reading interesting.

The word 'lying' seems a little uncomfortable when relating it to an alternative approach to finding the answer.

I agree about how "lying seems a little uncomfortable". It suggests that we're doing something wrong...

I'm pretty sure that's just personal moral attachments to the idea of "lying" and not anything implied by the text.

maybe "moral fudging"?

i might go with skillful fudging, or educated fudging.

$$\nabla \cdot \mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

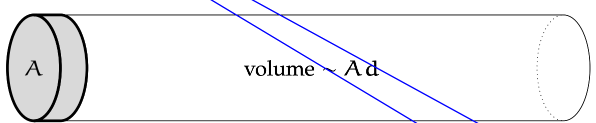

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

I don't really see how this involves symmetry. Unless your conception of symmetry is different than mine. It seems more like just simplification, I think that whatever link to symmetry that you see should be more explicitly explained.

I haven't finished reading, but I think the point it to see the invariant point (conservation of energy), and then use symmetry operations. (the roller-coaster example!).

I agree, I don't see how this is supposed to be "symmetry". All of the "symmetry" problems and readings we've had so far seem like arbitrary simplifications. Could we possibly have a more concrete definition of symmetry, and how it is applied?

I disagree that the simplifications were arbitrary - a lot of times, simplification can be done by looking at how to turn the problem you have into a problem that is similar but easier (ovals to circles, for example). It's all a matter of looking at a problem and thinking "oh, this would be sooo much easier if it looked like that other thing I know how to solve."

I think sometimes I get too caught up in the predetermined definition of "symmetry" in my head where it is linked to a visual image being symmetric about an axis. Calling it "invariant" clicks better for me.

I agree–if you think of it as "invariant", it makes a lot more sense–energy is conserved so if we lose some in one place, we gain it somewhere else. It's not symmetry in that geometric sense of a square having two symmetrical sides each with 2 sides at 90degrees; it's more like if we change one thing we know what happens to other things (once we find what stays constant or invariant).

Technically speaking, Noether's theorem says that time invariance will correspond to some conserved quantity: energy.

can we always use skillful lying, just curious

typo

typo here, remove one 'analysis'

$$\nabla \cdot \mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

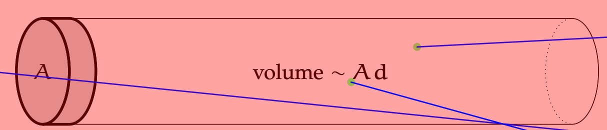

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

Besides the double typo, I'm not sure "The analysis imagines" makes any sense. Perhaps you mean something more like "For our analysis, imagine"?

Just a different way of saying something similar, I believe.

Is the double repetition a typo or is it intentional? Also, it is impossible to spell check your comments on NB. Normally on firefox, when typing in forms, a red squibble shows up underneath the word and you can right click to get a spelling suggestion. This does not work on NB though since right clicking opens up another comment box. Can you fix this?

I'm assuming this is a generalized case then? Since with the example of the cone, the velocity isn't constant. Or are we assuming it is?

Diagrams are always helpful. Having a picture which shows drag force in terms of Energy per distance greatly improves my understanding over the material.

I believe this example is being used to explain the drag force and therefore the problem is more simple to understand with a constant speed.

shouldn't it be faster because it is actually a cone rather than a cylinder?

My guess would also be that a cone is faster, but I think the basic concept is better taught using a cylinder as an example.

An arrow or something to show the path would be helpful. In this diagram I'm not sure if the puck is moving left to right or has already moved right to left.

By this figure if seems like the large cone would fall more slowly. I guess this also depends on how the surface area and mass scale.

Also you could separate the puck and the cylinder. That would clarify instantly that this diagram is an object trying to move through fluid, where as right now glancing at it, it just seems like a cap of a cylinder.

what do you mean by energy consumed?

$\nabla \cdot \mathbf{v} = 0.$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

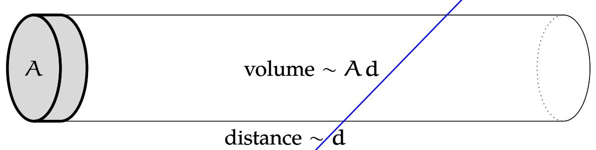

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

Is this supposed to be obvious? If so, it's not obvious to me. How do you arrive at this?

W=F*dx, or F=W/dx. The force, F, is the drag force, and W is the work done, or energy consumed. 'dx' is your unit distance.

For me this would have been helpful as part of an explicit definition of drag towards the start of the section.

I think it would be useful to have some other equations here to refresh those of us not well-read in the fields of physics.

More and more demands seem to be answered if the reader only continues reading before he asks them. How about reading an extra page and then asking for definitions and equations?

I think it be a poor design for a text to introduce a new and unfamiliar topic, only to define it pages later. It makes sense for people to request definitions and equations for drag, after now reading 2 pages about it and potentially being lost and confused. The non-course 6 people were confused and demanded explanations about UNIX, and now the non-course 2 people are confused and would like explanations about drag, and I think everyone was completely right to do so.

Really? I found this sentence to have the purpose of defining the drag force as energy consumed over distance. That's what it is.

How does viscosity remove KE?

I'm not sure about the exact details, but I can give a general explanation. Viscosity is the measure of the resistance of a fluid. Simply in one word, it is "thickness." So the thicker the fluid is, the more that the kinetic energy of the object decreases as it is traveling through the fluid.

$$\nabla \cdot \mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

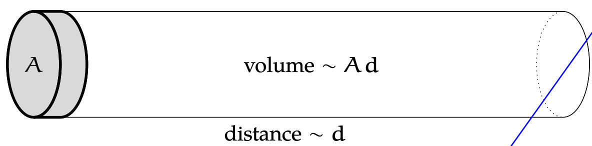

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

Kinetic energy is usually 1/2 times this quantity. Is this just for approximations sake?

Good question, I was wondering about this too. I assume it's just to make approximations easier since it really only affects the answer by a magnitude of 1 (half of 100 (10^2) is 50 (5*10^1))

It's because we don't care about constants, regardless of how large they are (in fact, I believe there are other constants missing, like the drag coefficient, which is geometry-dependent). What's important is how the force, below, scales with these constants. So really we have KE   M*v^2, M   rho*A*d, so KE   rho*A*v^2*d

That's why we're writing these with   instead of =, since it's rough estimation.

No, I believe   stands for "goes like" or "is proportional to" or "scales like" (take your pick). It is different, and in many ways more technical than "approximately equal to", which would be the 'double squiggle'.

The reason   is used, is because several constant factors are omitted. Since we only care about how two velocities relate to one another, they are ok to leave out. If you wanted to find out what the actual terminal velocity was, however, you'd have to include them (and determine their values, of course).

I think this would have been easier to understand written out as equations rather than into a paragraph. It is hard to follow math in word form.

Yeah: the   shows proportionality, since we're not including constants here. Remember, in the problem we're looking at how the two cones will fall; we don't care about their absolute velocities or drag, we only care about their velocities relative to each other. So we don't need to worry about constants (which will cancel out when we find the ratios anyways)

What is mass disturbed?

The mass of the volume displaced by the object (density*volume=mass)

$$\nabla \cdot \mathbf{v} = 0.$$

One imposes boundary conditions, which include the motion of the object and the requirement that no fluid enters the object – and solves for the pressure p and the velocity gradient at the surface of the object. Integrating the pressure force and the shear force gives the drag force.

In short, solving the equations analytically is difficult. I could spend hundreds of pages describing the mathematics to solve them. Even then, solutions are known only in a few circumstances, for example a sphere or a cylinder moving slowly in a viscous fluid or a sphere moving at any speed in an zero-viscosity fluid. But an inviscid fluid – what Feynman calls 'dry water' [9, Chapter II-40] – is particularly irrelevant to real life since viscosity is the reason for drag, so an inviscid solution predicts zero drag! Conservation of energy, supplemented with skillful lying, is a simple and quick alternative.

The analysis analysis imagines an object of cross-sectional area A moving through a fluid at speed v for a distance d:

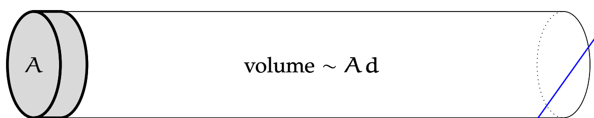

The drag force is the energy consumed per distance. The energy is consumed by imparting kinetic energy to the fluid, which viscosity eventually removes from the fluid. The kinetic energy is mass times velocity squared. The mass disturbed is $\rho A d$, where ρ is the fluid density (here, the air density). The velocity imparted to the fluid is roughly the velocity of the disturbance, which is v. So the kinetic energy imparted to the fluid is $\rho A v^2 d$, making the drag force

$$F \sim \rho A v^2.$$

The analysis has a divide-and-conquer tree:

I am continually impressed that all this analysis takes is a little concentrated thinking and application of easy 8.01 principles. These are such good exercises!

This is amazing. I admit, I always wondered where the v^2 term in drag force came from.

I agree with the above. This short paragraph is a great breakdown, and this type of explanation makes something very complex understandable for a student in high school physics.

Although I agree this is a clever, simple way to solve for the drag force, I don't think it's so useful in practice. The point of calculating drag forces (the "real" way) is to predict what final velocities will be or what forces on necessary to account for the lost energy. It's rare that you care about the drag force if you already know parameters like final velocity or energy lost.

Is the energy lost due to the fluid or the interaction between the fluid and the member it is traveling through?

This is mass disturbed per unit time right?

This is an amazing simplification of the previous set of equations

should probably define what these are, even if they seem obvious. the density threw me off, I wasn't sure which density you were referring too (fluid vs air?)

I really like how this book builds upon techniques from previous chapters. It's nice to be able to relate something new and unfamiliar (symmetry) to something we already understand (divide and conquer).

Agreed, I like the tie in with earlier sections. Although, to be honest, I'm having a tough time distinguishing between symmetry and D&C in some cases. What are the main differences?

Yeah I agree with both points. It is really nice to see that things we learn in the beginning of the class don't prove to be useless by the end of the class. I am also a little unclear on what symmetry is.

i can't describe how happy i was to see this...the above wouldn't have made nearly as much sense otherwise

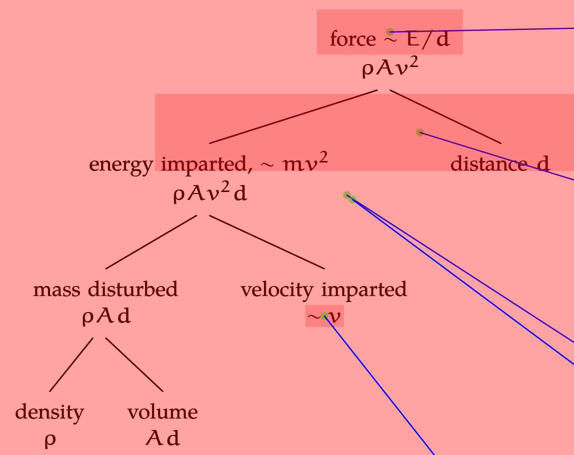

The result that $F_{drag} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly (see Problem 8.6), so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

$$v \sim \sqrt{\frac{F_{drag}}{\rho A}}.$$

The cross-sectional areas are easy to measure with a ruler, and the ratio between the small- and large-cone terminal velocities is even easier. The experiment is set up to make the drag force easy to measure: Since the cones fall at their respective terminal velocities, the drag force equals the weight. So

$$v \sim \sqrt{\frac{W}{\rho A}}.$$

Each cone's weight is proportional to its cross-sectional area, because they are geometrically similar and made out of the same piece of paper. So the terminal velocity v is independent of the area A: so the small and large cones should fall at the same speed.

To test this prediction, I stood on a table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another. So, the approximate conservation-of-energy analysis gains in plausibility (all the inaccuracies are hidden within the changing drag coefficient).

It would be nice if F_drag were at the very top of this tree (above force)

Building up a tree like this came earlier in my thought process about this problem, so could it go in earlier? Or at least the root could and we could start building branches as the reasoning progressed?

The "force" mentioned is the same as the "F_drag" in this problem. The idea is that we used energy and distance to determine the force, if we didn't know the drag force equation (which explains why we simplified the problem by removing certain coefficients).

Why didn't you use the 1 and -1 powers along the lines here?

Maybe he hasn't had time to edit this since he introduced those?

Also if he's using the tree generating code, maybe said code does not (yet?) have support for powers.

Vaguely related (but not entirely), the other day a friend showed me some tree-generation software called graphviz for those of us who want to make nice trees but don't know how to program (www.graphviz.org)

Like the use of earlier notes on diagrams

This does make everything so much easier

thank you!

I really like how the tree diagram is shown after the explanation of how we got the approximation for the force of drag. It makes a lot of sense and makes stuff more clear.

I think for all of your diagrams, figure captions would help a lot.

I think this diagram is self-explanatory. It actually helps a lot as it is...

I don't think a tree diagram needs any more figure captions than the info that's already on it. From here it's clear what breaks down into what

I agree. I get 80% from reading the above paragraph but the tree diagram cements it for me.

this is "approx" v right? what's the need of the "approx" if we're approximating anyway

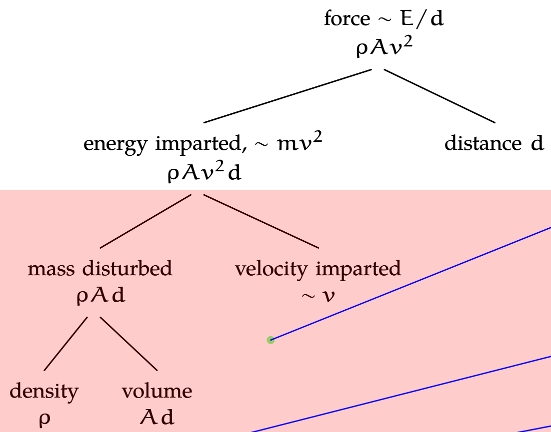

The result that $F_{\text{drag}} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly (see Problem 8.6), so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

$$v \sim \sqrt{\frac{F_{\text{drag}}}{\rho A}}.$$

The cross-sectional areas are easy to measure with a ruler, and the ratio between the small- and large-cone terminal velocities is even easier. The experiment is set up to make the drag force easy to measure: Since the cones fall at their respective terminal velocities, the drag force equals the weight. So

$$v \sim \sqrt{\frac{W}{\rho A}}.$$

Each cone's weight is proportional to its cross-sectional area, because they are geometrically similar and made out of the same piece of paper. So the terminal velocity v is independent of the area A: so the small and large cones should fall at the same speed.

To test this prediction, I stood on a table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another. So, the approximate conservation-of-energy analysis gains in plausibility (all the inaccuracies are hidden within the changing drag coefficient).

All of this could have been done with the powers along the lines instead of writing out the equations. It's nice to see them in writing, but I think you should be consistent throughout the book.

I also like the power method for it's simplicity and ease of reading. It would be nice to have the equations after though.

This is just like the coefficient of drag approach that I was talking about earlier. However, this is a more rudimentary and straightforward approach and is easier to derive/remember.

It might be better if the form of this was consistent with the rest of the document, so that it reads pAv^2.

Agreed.

So can I assume that the p=1 if it's air? In which case it all comes down to the Av^2?

This still seems like a lot of work for two cones that will hit the ground at nearly the same time

It seems that because the result is simple. But the method is a lot less work than solving the Navier-Stokes equations to get the same result.

Also, lots of other results are plausible: for example that the bigger cone is twice as fast. For highly viscous flow (e.g. big and small cones in honey), it turns out that the two objects do not hit at the same time. That analysis may end up on a problem set...

That would be a cool follow-up.

I don't remember the distribution, but the first day of class not everyone got the right answer. This isn't an entirely intuitive answer.

Where is this problem? I couldn't find it.

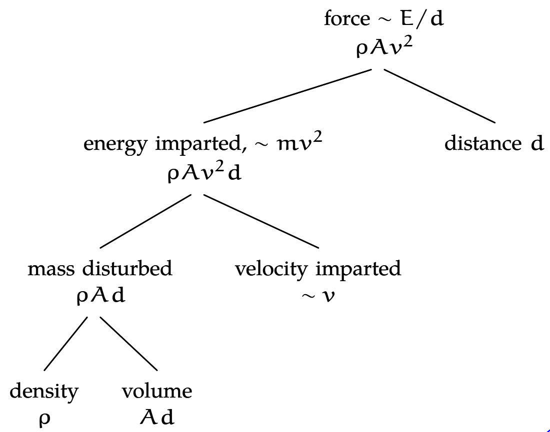

The result that $F_{drag} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly (see Problem 8.6), so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

$$v \sim \sqrt{\frac{F_{drag}}{\rho A}}.$$

The cross-sectional areas are easy to measure with a ruler, and the ratio between the small- and large-cone terminal velocities is even easier. The experiment is set up to make the drag force easy to measure: Since the cones fall at their respective terminal velocities, the drag force equals the weight. So

$$v \sim \sqrt{\frac{W}{\rho A}}.$$

Each cone's weight is proportional to its cross-sectional area, because they are geometrically similar and made out of the same piece of paper. So the terminal velocity v is independent of the area A: so the small and large cones should fall at the same speed.

To test this prediction, I stood on a table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another. So, the approximate conservation-of-energy analysis gains in plausibility (all the inaccuracies are hidden within the changing drag coefficient).

This makes sense intuitively (it is not difficult to imagine the cones not accelerating too much after a second or two), but is there another way, perhaps more rigorous, to come to this conclusion?

You can estimate a drag coefficient, find a more precise equation for v, and get an estimate of terminal velocity. From there, you can take a stab at figuring out its acceleration (which will depend on the cone radius), and then estimate how long it will take to reach terminal velocity.

Is there an abstraction here? As in some time constant that can represent how quickly something goes from rest to terminal velocity?

I think if you just think about the quantities involved it falls out pretty naturally. Drag should be proportional to cross sectional area and the weight is proportional to the mass. So for paper cones which have a "large" cross sectional area and a "small" mass you would expect the forces to balance out pretty quickly.

I think it might be useful to first put newton's second law and equate the forces before you skip straight to the terminal velocity equation. Just so that people know where this terminal velocity term actually comes from

Is this always the _terminal_ velocity? Above when we found the relation of Fdrag to pAv^2, I didn't think we had specified that.

The definition of terminal velocity is defined when these quantities are equal, or when the drag forces are equal to the opposing forces whether it be cycling or falling. This was intuitive to me, but it may help to define the terminal velocity for those that haven't seen a problem like this before.

I think you may have missed 2:55's point. The equation that immediately follows this use of "terminal velocity" is just a rearrangement of the drag equation (F_drag pAv^2). The velocity the drag equation uses is just the velocity of the object with respect to the fluid.

Terminal velocity occurs once we've replaced Fdrag with Weight (since at that point there's no acceleration and thus no change in velocity).

So I think this use of the phrase "terminal velocity" is mistaken.

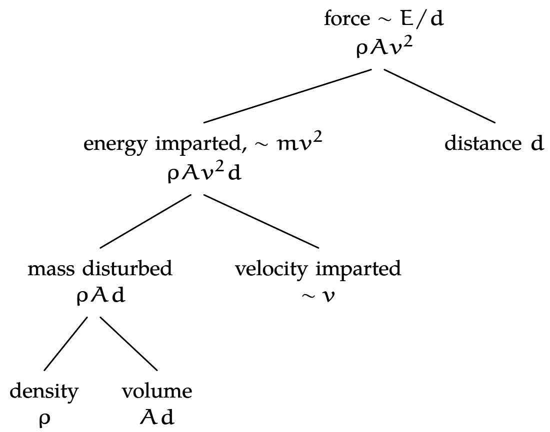

The result that $F_{drag} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly (see Problem 8.6), so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

$$v \sim \sqrt{\frac{F_{drag}}{\rho A}}.$$

The cross-sectional areas are easy to measure with a ruler, and the ratio between the small- and large-cone terminal velocities is even easier. The experiment is set up to make the drag force easy to measure: Since the cones fall at their respective terminal velocities, the drag force equals the weight. So

$$v \sim \sqrt{\frac{W}{\rho A}}.$$

Each cone's weight is proportional to its cross-sectional area, because they are geometrically similar and made out of the same piece of paper. So the terminal velocity v is independent of the area A: so the small and large cones should fall at the same speed.

To test this prediction, I stood on a table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another. So, the approximate conservation-of-energy analysis gains in plausibility (all the inaccuracies are hidden within the changing drag coefficient).

I feel like a good portion of readers wont actually try this experiment... i think it would help to have a quick sentence concluding the findings so I'm not as uncertain?

He does summarize his findings a little later in the reading.

Good idea for an introduction to the subject!

This is an interesting thought but why not make this inference in the beginning of the problem.

Why W instead of mg? I feel like the concept that acceleration is 0 is left out here when that is what tells us the drag and weight forces cancel. Free body diagram?

I wouldn't ponder too much over asking why W instead of mg, but I would defnitely agree that a free-body diagram would be extremely useful here.

I just think "work" when I see W (as someone used it in earlier comments), instead of Fw or mg.

so it turns out is does depend on the ratio between surface area and weight

It might be useful here to write down in equation form that the weight and cross-sectional area are proportional to each other by a factor of 1, which is why the terminal velocity is independent of area.

Agreed.

wait....how are they proportional by a factor of 1? Maybe im just tired but I dont see this.

I don't get this transition. I understand that each cone's weight is proportional to its area by why does this mean that mean the velocity is independent of area?

So are we saying that for this case (same shape & same material), if you triple the area you also triple the weight? Thus the velocity will be the same no matter which 3 variables we change.

So are we saying that for this case (same shape & same material), if you triple the area you also triple the weight? Thus the velocity will be the same no matter which 3 variables we change.

I'm happy that we saw this demonstration in class! It makes the example so much easier to relate to!

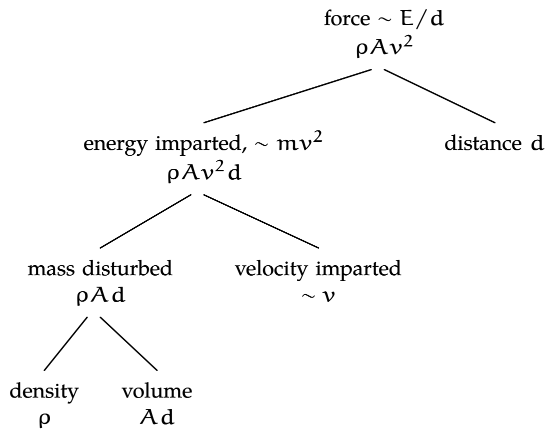

The result that $F_{drag} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly (see Problem 8.6), so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

$$v \sim \sqrt{\frac{F_{drag}}{\rho A}}.$$

The cross-sectional areas are easy to measure with a ruler, and the ratio between the small- and large-cone terminal velocities is even easier. The experiment is set up to make the drag force easy to measure: Since the cones fall at their respective terminal velocities, the drag force equals the weight. So

$$v \sim \sqrt{\frac{W}{\rho A}}.$$

Each cone's weight is proportional to its cross-sectional area, because they are geometrically similar and made out of the same piece of paper. So the terminal velocity v is independent of the area A: so the small and large cones should fall at the same speed.

To test this prediction, I stood on a table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another. So, the approximate conservation-of-energy analysis gains in plausibility (all the inaccuracies are hidden within the changing drag coefficient).

I found the demonstration in class helpful.

but this might not be so accurate though because your arms might be at different heights, but then again, this is an approximation class

When you performed the experiment in class, it actually looked like the cone fell at different speeds, but I guess it's difficult to be precise.

Does this mean that drag doesn't have as much effect as we thought it would?.. or that weight and drag scale similarly?

It would be helpful to see the numbers from the expirement plugged into the equations to obtain this result.

What would be interesting would be to predict this time difference. It should be integral(Vsmall(t)-Vbig(t))/Vterminal. Basically we'd have to estimate the area between the velocity curves.

just as we saw in the experiment...yippie! :)

how are the inaccuracies hidden? I understand that we significantly simplified the problem and will therefore have errors, perhaps these errors should be explained in more detail?

so the drag coefficient is not constant, but we assume it is constant?

Are they hidden or we just don't consider them? Are we deeming the inaccuracies insignificant, or are they really not there...?

I think he means hidden away as in we used abstraction to make the problem a much simpler one–remember all those crazy equations in the beginning

Again, the idea of abstraction comes into play as well here. We have abstracted away all the nasty math into one variable - drag coefficient.

That abstraction was amazing for understanding - the drag coefficient in a small equation is much easier to understand than a whole mess of math that I never want to see again.

These three pages are really good, minus the one typo. I think the explanation is really clean and very easy to follow. If there is ever a question on how a subject should be brought up, this would be a great template.

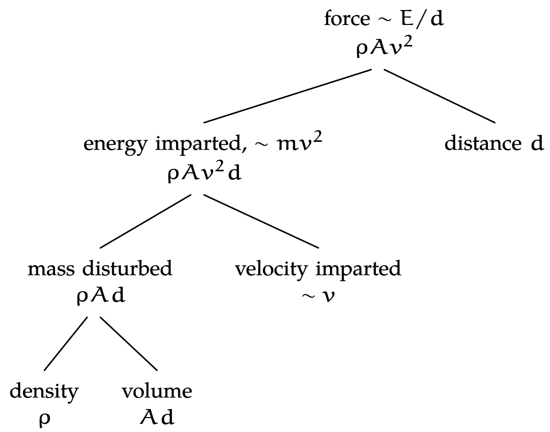

The result that $F_{drag} \sim \rho v^2 A$ is enough to predict the result of the cone experiment. The cones reach terminal velocity quickly (see Problem 8.6), so the relevant quantity in finding the fall time is the terminal velocity. From the drag-force formula, the terminal velocity is

$$v \sim \sqrt{\frac{F_{drag}}{\rho A}}.$$

The cross-sectional areas are easy to measure with a ruler, and the ratio between the small- and large-cone terminal velocities is even easier. The experiment is set up to make the drag force easy to measure: Since the cones fall at their respective terminal velocities, the drag force equals the weight. So

$$v \sim \sqrt{\frac{W}{\rho A}}.$$

Each cone's weight is proportional to its cross-sectional area, because they are geometrically similar and made out of the same piece of paper. So the terminal velocity v is independent of the area A: so the small and large cones should fall at the same speed.

To test this prediction, I stood on a table and dropped the two cones. The fall lasted about two seconds, and they landed within 0.1 s of one another. So, the approximate conservation-of-energy analysis gains in plausibility (all the inaccuracies are hidden within the changing drag coefficient).

wait, what did this have to do with symmetry again?

When you did this in class, I figured they'd fall around the same time without equations or math – just figuring that the increased surface area and weight canceled each other out. Does that count as an estimation method?

I had the same intuition- I think that method of estimating is more like the "gut feeling" that Sanjoy has talked about. This method is better for estimation because it gives us concrete evidence for our estimation as opposed to a general feeling.

This example is one where thinking about it a little bit often leads to the right answer, as does thinking about it a lot (e.g. the analysis in the text).

But thinking about it a moderate amount can lead to puzzlement. For example, should the drag be proportional to area or to radius? For very slow or viscous flows (e.g. a fog droplet), as you will learn when we do extreme-cases reasoning, the drag is proportional to radius.

Therefore, maybe the experiment to introduce drag is four small cones versus one small one, which I'll do on Wednesday in lecture.

I think this gut decision is hard to make - without knowledge of what ratios and exponentials are present in the equations, we cannot determine whether weight, surface area, or neither dominates the situation.

3.4 Cycling

This section discusses cycling as an example of how drag affects the performance of people as well as fleas. Those results will be used in the analysis of swimming, the example of the next section.

What is the world-record cycling speed? Before looking it up, predict it using armchair proportional reasoning. The first task is to define the kind of world record. Let's say that the cycling is on a level ground using a regular bicycle, although faster speeds are possible using special bicycles or going downhill.

To estimate the speed, make a model of where the energy goes. It goes into rolling resistance, into friction in the chain and gears, and into drag. At low speeds, the rolling resistance and chain friction are probably important. But the importance of drag rises rapidly with speed, so at high-enough speeds, drag is the dominant consumer of energy.

For simplicity, assume that drag is the only consumer of energy. The maximum speed happens when the power supplied by the rider equals the power consumed by drag. The problem therefore divides into two estimates: the power consumed by drag and the power that an athlete can supply.

The drag power P_{drag} is related to the drag force:

$$P_{drag} = F_{drag}v \sim \rho v^3 A.$$

It indeed rises rapidly with velocity, supporting the initial assumption that drag is the important effect at world-record speeds.

Setting $P_{drag} = P_{athlete}$ gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}$$

To estimate how much power an athlete can supply, I ran up one flight of stairs leading from the MIT Infinite Corridor. The Infinite Corridor, being an old building, has spacious high ceilings, so the vertical climb is perhaps $h \sim 4$ m (a typical house is 3 m per storey). Leaping up the stairs as fast as I could, I needed $t \sim 5$ s for the climb. My mass is 60 kg, so my power output was

Is this supposed to be due Thursday?

Tuesday (with Sec 3.3). For Thursday you'll learn about flight.

Cool!

This phrasing seems a bit awkward - the meaning is clear, but it just bothers me a bit.

non sequitur

I had to read it a few times before I understood this sentence.

Are we going to talk about fleas in later sections, or is this just to speak to the ubiquity of this content?

Why fleas?? haha

These results? As in, the results of this section?

Yeah, this doesn't quite make sense in the context.

speed skating or skiing would be more relevant right now.

Maybe true but winter olympics only come around every 4 years...

but he's hoping to write a textbook! hopefully it will be read by many people for many olympics to come :)

Simply due to the Olympics?

I think that, in general, more people are able to relate to swimming over speed skating or skiing...

I think the first comment was a just a joke. Swimming is a perfectly relevant example.

I would say more people have direct experience swimming than speed skating or skiing... especially since we all have to pass the swim test :)

SWIMMING FTW

100 MPH?

3.4 Cycling

This section discusses cycling as an example of how drag affects the performance of people as well as fleas. Those results will be used in the analysis of swimming, the example of the next section.

What is the world-record cycling speed? Before looking it up, predict it using armchair proportional reasoning. The first task is to define the kind of world record. Let's say that the cycling is on a level ground using a regular bicycle, although faster speeds are possible using special bicycles or going downhill.

To estimate the speed, make a model of where the energy goes. It goes into rolling resistance, into friction in the chain and gears, and into drag. At low speeds, the rolling resistance and chain friction are probably important. But the importance of drag rises rapidly with speed, so at high-enough speeds, drag is the dominant consumer of energy.

For simplicity, assume that drag is the only consumer of energy. The maximum speed happens when the power supplied by the rider equals the power consumed by drag. The problem therefore divides into two estimates: the power consumed by drag and the power that an athlete can supply.

The drag power P_{drag} is related to the drag force:

$$P_{drag} = F_{drag} v \sim \rho v^3 A.$$

It indeed rises rapidly with velocity, supporting the initial assumption that drag is the important effect at world-record speeds.

Setting $P_{drag} = P_{athlete}$ gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}$$

To estimate how much power an athlete can supply, I ran up one flight of stairs leading from the MIT Infinite Corridor. The Infinite Corridor, being an old building, has spacious high ceilings, so the vertical climb is perhaps $h \sim 4\,m$ (a typical house is 3 m per storey). Leaping up the stairs as fast as I could, I needed $t \sim 5\,s$ for the climb. My mass is 60 kg, so my power output was

is there a technical term for this? i feel like there is... i just don't know it off the top of my head...

Well, the techinical definition isn't so important. Really, we just need to make it clear so that if we compare answers or methods, we don't have conflicts based on having calculated multiple scenarios. This is a good point - I would have certainly missed it.

Armchair proportional reasoning? I'm unfamiliar with this term.

I think this means just sit around and think about it (as though you were in a comfy armchair) rather than look it up

Yeah, I think this is an awkward term and could be rephrased better.

The first thing that came to my mind when I read this was that we should approximate a cyclist as having the cross sectional area of an armchair...

Yeah, same. I tried to think of ways approximating an armchair would actually be helpful.

This makes me think that you are referring to the instantaneous speed record, which is quite different from the one hour average speed you used in the example, it should be clear before hand what exactly you are trying to approximate.

I agree. When I first read this I was trying to think of "whats the fastest anyone has ever pedaled a bicycle? were they on a hill? etc."

"Using a regular bicycle" is a very dangerous term. A mountain bike's max speed will differ from that of a road bike, which will differ from that of an expensive road bike, which will differ from that of a time trial bike, which will differ from that of a triathlon bike....

These differences are also quite considerable, even just between types and models of road bike. Indeed, even the helmet used can effect the speed by a decent amount.

3.4 Cycling

This section discusses cycling as an example of how drag affects the performance of people as well as fleas. Those results will be used in the analysis of swimming, the example of the next section.

What is the world-record cycling speed? Before looking it up, predict it using armchair proportional reasoning. The first task is to define the kind of world record. Let's say that the cycling is on a level ground using a regular bicycle, although faster speeds are possible using special bicycles or going downhill.

To estimate the speed, make a model of where the energy goes. It goes into rolling resistance, into friction in the chain and gears, and into drag. At low speeds, the rolling resistance and chain friction are probably important. But the importance of drag rises rapidly with speed, so at high-enough speeds, drag is the dominant consumer of energy.

For simplicity, assume that drag is the only consumer of energy. The maximum speed happens when the power supplied by the rider equals the power consumed by drag. The problem therefore divides into two estimates: the power consumed by drag and the power that an athlete can supply.

The drag power P_{drag} is related to the drag force:

$$P_{drag} = F_{drag}v \sim \rho v^3 A.$$

It indeed rises rapidly with velocity, supporting the initial assumption that drag is the important effect at world-record speeds.

Setting $P_{drag} = P_{athlete}$ gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}$$

To estimate how much power an athlete can supply, I ran up one flight of stairs leading from the MIT Infinite Corridor. The Infinite Corridor, being an old building, has spacious high ceilings, so the vertical climb is perhaps $h \sim 4$ m (a typical house is 3 m per storey). Leaping up the stairs as fast as I could, I needed $t \sim 5$ s for the climb. My mass is 60 kg, so my power output was

Being a racing cyclist, I cannot agree more with your choice of example and the analysis here. Air drag is so incredibly annoying. So much of competitive cycling involves trying to find the tiniest changes in equipment and rider form to reduce air drag. It can get ridiculous pretty fast, in fact there are even men who shave their legs becuase they think it will cause less air drag and make them go faster.

Figuring out whether shaving their legs would cause less air drag would be an interesting problem.

Listing all the ways energy can get dissipated in riding a bicycle is a good way to see the problem. Could we also be able to do a free body diagram while performing a force balance order to find the drag force? (since acceleration=0 at terminal velocity)

In the D-Lab wheelchair design class, we do the opposite calculations...when determining the power output of a person, you ignore drag completely because in a wheelchair you never travel fast enough for drag to matter more than 10%. Interesting to see the other end of the spectrum here.

What about motorized wheelchairs? Doesn't drag come into play? Also, what about windy days? The cross section of someone on a wheelchair seems to be bigger.

Is this necessary? Couldn't we just estimate it based on other types of vehicles and based on power?

Power comes into the analysis later on. We could use other vehicles, but it might take longer time since we need to compare the fuel used by the vehicle and the energy/strength of a person.

I like that in the previous section we saw the equation for drag, so we are familiar here with the fact that drag and velocity are directly proportional.

I agree, using the same concepts twice is really useful for understanding the estimations and not getting hooked up on the equations.

Yeah I agree too. Seeing as I am not very familiar with the concept of drag, it was very helpful to have seen the simple cone example before reading this section.

But it's not a linear proportion (see below). This is why driving your car at 65 mph burns nonlinearly more gas than at 55.

So do you just not think about friction in this case?

3.4 Cycling

This section discusses cycling as an example of how drag affects the performance of people as well as fleas. Those results will be used in the analysis of swimming, the example of the next section.

What is the world-record cycling speed? Before looking it up, predict it using armchair proportional reasoning. The first task is to define the kind of world record. Let's say that the cycling is on a level ground using a regular bicycle, although faster speeds are possible using special bicycles or going downhill.

To estimate the speed, make a model of where the energy goes. It goes into rolling resistance, into friction in the chain and gears, and into drag. At low speeds, the rolling resistance and chain friction are probably important. But the importance of drag rises rapidly with speed, so at high-enough speeds, drag is the dominant consumer of energy.

For simplicity, assume that drag is the only consumer of energy. The maximum speed happens when the power supplied by the rider equals the power consumed by drag. The problem therefore divides into two estimates: the power consumed by drag and the power that an athlete can supply.

The drag power P_{drag} is related to the drag force:

$$P_{drag} = F_{drag}v \sim \rho v^3 A.$$

It indeed rises rapidly with velocity, supporting the initial assumption that drag is the important effect at world-record speeds.

Setting $P_{drag} = P_{athlete}$ gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}$$

To estimate how much power an athlete can supply, I ran up one flight of stairs leading from the MIT Infinite Corridor. The Infinite Corridor, being an old building, has spacious high ceilings, so the vertical climb is perhaps $h \sim 4\,m$ (a typical house is 3 m per storey). Leaping up the stairs as fast as I could, I needed $t \sim 5\,s$ for the climb. My mass is 60 kg, so my power output was

I feel like the (from my feelings when riding a bike) rolling resistance should be rather important...even at high speeds.

This is a really cool way to solve the problem

Ah, perhaps I could have used this for the last problem set asking about the force of wind on a cyclist!

Indeed, and you still can use it. That's why I put the problem on HW 03.

This does seem useful for the current problem set - I had a tough time figuring out how to approach that one.

Do the mathematics for calculating drag have any relation to calculating draft...like drafting off of a car or a cyclist etc.??

I believe drafting is a much more similar mathematical calculation but relies on the same principal. The person that is drafting is at an advantage because they don't feel the drag.

We can see from this section that drag is the main force of resistance to any moving object, so if we can eliminate it by drafting we save a surprisingly (or not) huge amount of energy

It's spelled story in the U.S. and storey in Britain. I'm not sure which is the principal audience...

Interesting way to approximate this.

Haha, agreed. I would have definitely used some other sort of method, but actual, physical experimentation is also nice! I sometimes forget about that as we delve further into thought experiments in this class.

If you don't have a dynamometer, this is a good substitute.

My physics teacher in high school used this example a lot. He talked about how much power was exerted running up all the stairs in the green building (he went to MIT).

3.4 Cycling

This section discusses cycling as an example of how drag affects the performance of people as well as fleas. Those results will be used in the analysis of swimming, the example of the next section.

What is the world-record cycling speed? Before looking it up, predict it using armchair proportional reasoning. The first task is to define the kind of world record. Let's say that the cycling is on a level ground using a regular bicycle, although faster speeds are possible using special bicycles or going downhill.

To estimate the speed, make a model of where the energy goes. It goes into rolling resistance, into friction in the chain and gears, and into drag. At low speeds, the rolling resistance and chain friction are probably important. But the importance of drag rises rapidly with speed, so at high-enough speeds, drag is the dominant consumer of energy.

For simplicity, assume that drag is the only consumer of energy. The maximum speed happens when the power supplied by the rider equals the power consumed by drag. The problem therefore divides into two estimates: the power consumed by drag and the power that an athlete can supply.

The drag power P_{drag} is related to the drag force:

$$P_{drag} = F_{drag}v \sim \rho v^3 A.$$

It indeed rises rapidly with velocity, supporting the initial assumption that drag is the important effect at world-record speeds.

Setting $P_{drag} = P_{athlete}$ gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}$$

To estimate how much power an athlete can supply, I ran up one flight of stairs leading from the MIT Infinite Corridor. The Infinite Corridor, being an old building, has spacious high ceilings, so the vertical climb is perhaps $h \sim 4$ m (a typical house is 3 m per storey). Leaping up the stairs as fast as I could, I needed $t \sim 5$ s for the climb. My mass is 60 kg, so my power output was

Brilliant!

This is really cool...as is the explanation of the estimation for olympic athletes in comparison

I agree–this way of finding out power output is so easy that we (I do at least) overlook it–just think of real life examples, or better yet, do them!

$$P_{author} \sim \frac{\text{potential energy supplied}}{\text{time to deliver it}}$$
$$= \frac{mgh}{t} \sim \frac{60\,\text{kg} \times 10\,\text{m}\,\text{s}^{-2} \times 4\,\text{m}}{5\,\text{s}} \sim 500\,\text{W}.$$

$P_{athlete}$ should be higher than this peak power since most authors are not Olympic athletes. Fortunately I'd like to predict the endurance record. An Olympic athlete's long-term power might well be comparable to my peak power. So I use $P_{athlete} = 500\,\text{W}$.

The remaining item is the cyclist's cross-sectional area A. Divide the area into width and height. The width is a body width, perhaps 0.4 m. A racing cyclist crouches, so the height is maybe 1 m rather than a full 2 m. So $A \sim 0.4\,\text{m}^2$.

Here is the tree that represents this analysis:

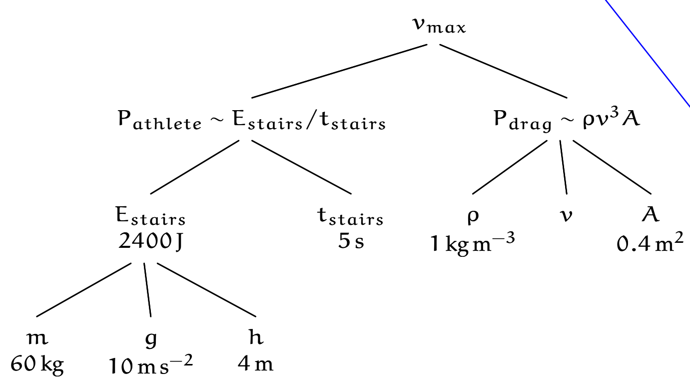

Now combine the estimates to find the maximum speed. Putting in numbers gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3} \sim \left(\frac{500\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3}.$$

The cube root might suggest using a calculator. However, massaging the numbers simplifies the arithmetic enough to do it mentally. If only the power were 400 W or, instead, if the area were 0.5 m! Therefore, in the words of Captain Jean-Luc Picard, 'make it so'. The cube root becomes easy:

$$v_{max} \sim\sim \left(\frac{400\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3} \sim (1000\,\text{m}^3\,\text{s}^{-3})^{1/3} = 10\,\text{m}\,\text{s}^{-1}.$$

Comments on page 5

Can we use these methods during the final? I think that'd be wicked funny.

To me, this seems way too specific for one person. Some people would do the same thing and go twice as fast, or twice as slow. I realize factors of two aren't that important in this class, but when trying to determine a world record speed, a factor of two seems too large to even be worth approximating.

But, importantly, Power is raised to the power of (1/3) in the above equation. So even a factor of 2 becomes just a factor of 1.26.

If it were, on the other hand, raised to the power of 3, then that factor of 2 would be amplified and could end up being about a factor of 8 in our final result.

it might be useful to compare this number to a light bulb or some other familiar source of power to get an idea of how useful we are as sources.

This wasn't really clear in the introduction.. I thought you just meant world record top speed.

I agree, a bit more definition in the beginning would help.

Yes, this can be stated here also (for the estimation effect), but needs to be in the intro to cycling.

What do you mean by endurance record...? As in the power output the athlete has on average?

I thought that was a joke, as we just needed a ballpark number to compare.

this is a little hard to follow

cycling and running up the stairs deliver power in different ways. is it prudent to use the same numbers?

I like this estimate

Me too.

I think it's fine but thinking about cyclists, I'm sure they're power output is almost 1.5x that of normal humans.

$$P_{author} \sim \frac{\text{potential energy supplied}}{\text{time to deliver it}}$$
$$= \frac{mgh}{t} \sim \frac{60\,\text{kg} \times 10\,\text{m}\,\text{s}^{-2} \times 4\,\text{m}}{5\,\text{s}} \sim 500\,\text{W}.$$

$P_{athlete}$ should be higher than this peak power since most authors are not Olympic athletes. Fortunately I'd like to predict the endurance record. An Olympic athlete's long-term power might well be comparable to my peak power. So I use $P_{athlete} = 500\,\text{W}$.

The remaining item is the cyclist's cross-sectional area A. Divide the area into width and height. The width is a body width, perhaps 0.4 m. A racing cyclist crouches, so the height is maybe 1 m rather than a full 2 m. So $A \sim 0.4\,\text{m}^2$.

Here is the tree that represents this analysis:

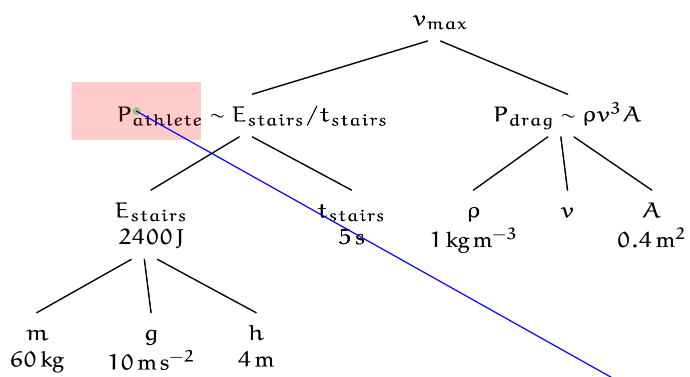

Now combine the estimates to find the maximum speed. Putting in numbers gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3} \sim \left(\frac{500\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3}.$$

The cube root might suggest using a calculator. However, massaging the numbers simplifies the arithmetic enough to do it mentally. If only the power were 400 W or, instead, if the area were 0.5 m! Therefore, in the words of Captain Jean-Luc Picard, 'make it so'. The cube root becomes easy:

$$v_{max} \sim\sim \left(\frac{400\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3} \sim (1000\,\text{m}^3\,\text{s}^{-3})^{1/3} = 10\,\text{m}\,\text{s}^{-1}.$$

interesting. why not 0.3 (few/10) or 0.5 (one-half)?

seeing the later calc, i understand why 0.4 makes things easier. is that why you choose it?

I'd like to know this as well - how would you have the foresight to estimate something as .4 instead of .few?

I also wonder. Plus, he later changed 500 to a 400 to match the 0.4, he could've started with a 0.5 and gotten the same thing.

I've always estimated the body to be 1m^2. But in this case the cyclist is crouching. Attention to detail got me on this one for the practice test.

This tree confuses me a little bit; it's different from the others we were using. Is a tree the best analysis tool here?

I think it would be helpful to see a tree before the analysis. It can have the values or not, but it helps me to understand why you are making certain calculations.

I agree. I also think a few of the previous problems we've seen in the reading for this course could use a bit of a roadmap (either in the form of a tree or written description) which could help indicate the assumptions/facts we need to focus on as we read through the example.

I like the tree analysis. With all the different components that combine to calculate the top speed, its nice to lay them out in a tree. It helps me reiterate my process and double check the work before diving into calculations.

Perhaps at the beginning, we could have a simple tree with vmax, Pathlete, and Pdrag. Then, when we start estimating Pathlete, we could draw that part of the tree, and then, once we've created all the subtrees, we can draw them together as one final tree.

This would have the advantage of giving us a general outline from the start without presenting confusing information (having 'Estairs' on the tree before it's mentioned in the text, for example, would just be confusing).

I'm not a fan of this font.

$$P_{\text{author}} \sim \frac{\text{potential energy supplied}}{\text{time to deliver it}}$$
$$= \frac{mgh}{t} \sim \frac{60\,\text{kg} \times 10\,\text{m}\,\text{s}^{-2} \times 4\,\text{m}}{5\,\text{s}} \sim 500\,\text{W}.$$

P_{athlete} should be higher than this peak power since most authors are not Olympic athletes. Fortunately I'd like to predict the endurance record. An Olympic athlete's long-term power might well be comparable to my peak power. So I use $P_{\text{athlete}} = 500\,\text{W}$.

The remaining item is the cyclist's cross-sectional area A. Divide the area into width and height. The width is a body width, perhaps 0.4 m. A racing cyclist crouches, so the height is maybe 1 m rather than a full 2 m. So $A \sim 0.4\,\text{m}^2$.

Here is the tree that represents this analysis:

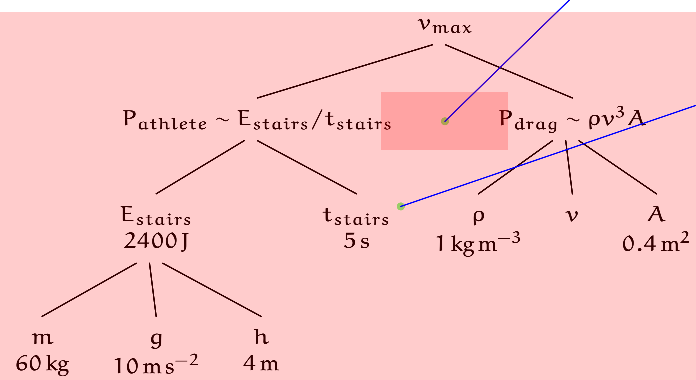

Now combine the estimates to find the maximum speed. Putting in numbers gives

$$v_{\max} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3} \sim \left(\frac{500\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3}.$$

The cube root might suggest using a calculator. However, massaging the numbers simplifies the arithmetic enough to do it mentally. If only the power were 400 W or, instead, if the area were 0.5 m! Therefore, in the words of Captain Jean-Luc Picard, 'make it so'. The cube root becomes easy:

$$v_{\max} \sim\sim \left(\frac{400\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3} \sim (1000\,\text{m}^3\,\text{s}^{-3})^{1/3} = 10\,\text{m}\,\text{s}^{-1}.$$

In class I think we used a dotted line to note that these values should be equal- it'd be nice to see that here.

We were also using it for redundancy, but these values are not redundant, they just come together to give us vmax.

Yes, I think overloading the dotted line notation might be confusing, but some sort of indication could be helpful.

yeah, in class a dotted line showed redundancy, but I dont think this is a case of redundancy...we're finding vmax based on Pdrag and Pathlete...

But based on the notation we learned this diagram says they should be multiplied, which is wrong. An improvement needs to be made to the diagram system that better explains the relationships between the nodes.

This tree is very helpful in explaining the analysis.

so in HW, we need to specify the range of error (or deviation), so if you just massage the numbers here, how do you account for it when you calculate your range of error

$$P_{\text{author}} \sim \frac{\text{potential energy supplied}}{\text{time to deliver it}}$$
$$= \frac{mgh}{t} \sim \frac{60\,\text{kg} \times 10\,\text{m}\,\text{s}^{-2} \times 4\,\text{m}}{5\,\text{s}} \sim 500\,\text{W}.$$

P_{athlete} should be higher than this peak power since most authors are not Olympic athletes. Fortunately I'd like to predict the endurance record. An Olympic athlete's long-term power might well be comparable to my peak power. So I use $P_{\text{athlete}} = 500\,\text{W}$.

The remaining item is the cyclist's cross-sectional area A. Divide the area into width and height. The width is a body width, perhaps 0.4 m. A racing cyclist crouches, so the height is maybe 1 m rather than a full 2 m. So $A \sim 0.4\,\text{m}^2$.

Here is the tree that represents this analysis:

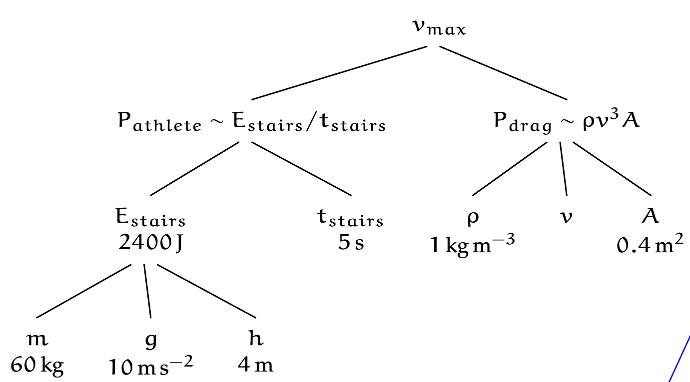

Now combine the estimates to find the maximum speed. Putting in numbers gives

$$v_{\max} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3} \sim \left(\frac{500\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3}.$$

The cube root might suggest using a calculator. However, massaging the numbers simplifies the arithmetic enough to do it mentally. If only the power were 400 W or, instead, if the area were 0.5 m! Therefore, in the words of Captain Jean-Luc Picard, 'make it so'. The cube root becomes easy:

$$v_{\max} \sim\sim \left(\frac{400\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3} \sim (1000\,\text{m}^3\,\text{s}^{-3})^{1/3} = 10\,\text{m}\,\text{s}^{-1}.$$

Is it really worth it to do it mentally? It's true that you get the answer a bit faster, but you might be sacrificing accuracy unnecessarily. I'm not quite sure how I feel about it, but it's something I thought of while reading the section.

I guess it depends on how much you care about accuracy. Since he's willing to fix the equation to suit his needs he must not care so much about being really close to the correct answer as he is of getting an answer that is on the order of the answer and that is reasonably close.

Is this really a question of whether or not to do it mentally? The answer in an Approximation class will always be to do it mentally. Since the beginning of class, we've been learning methods about how to make calculations simpler and easier to do 'mentally.'

Thats true, it does work reasonally well. On the other hand, I don't think I'd be confident enough to try this approach. If the number we were calculating is orders of magnitude higher, thats a different question. When calculating top speed, a few m/s off stands out, so I think I'd need some more practice with this.

I actually think we lose a lot these days because of the reliance on calculators. Yes, it does reduce some error, but I find my mental math to be greatly inferior to my mental math when I was in 4th grade...

nice, he is my favorite captain.

I agree that the Star Trek reference is awesome, but if this is a book for a general audience as opposed to just MIT, this may be lost on some people.

Was lost on me, but it could be intended for an MIT audience, or at least people who would be interested enough to look something like that up.

While the specific reference is lost on me, I like the point he's making with it. 'Make it so' is certainly useful.

Agreed. Maybe you may want to include a footnote defining who he is.

using the few method!

$$P_{author} \sim \frac{\text{potential energy supplied}}{\text{time to deliver it}}$$
$$= \frac{mgh}{t} \sim \frac{60\,\text{kg} \times 10\,\text{m}\,\text{s}^{-2} \times 4\,\text{m}}{5\,\text{s}} \sim 500\,\text{W}.$$

$P_{athlete}$ should be higher than this peak power since most authors are not Olympic athletes. Fortunately I'd like to predict the endurance record. An Olympic athlete's long-term power might well be comparable to my peak power. So I use $P_{athlete} = 500\,\text{W}$.

The remaining item is the cyclist's cross-sectional area A. Divide the area into width and height. The width is a body width, perhaps 0.4 m. A racing cyclist crouches, so the height is maybe 1 m rather than a full 2 m. So $A \sim 0.4\,\text{m}^2$.

Here is the tree that represents this analysis:

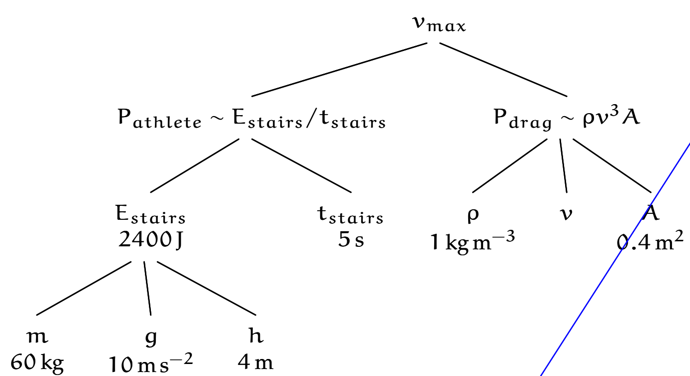

Now combine the estimates to find the maximum speed. Putting in numbers gives

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3} \sim \left(\frac{500\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3}.$$

The cube root might suggest using a calculator. However, massaging the numbers simplifies the arithmetic enough to do it mentally. If only the power were 400 W or, instead, if the area were 0.5 m! Therefore, in the words of Captain Jean-Luc Picard, 'make it so'. The cube root becomes easy:

$$v_{max} \sim\sim \left(\frac{400\,\text{W}}{1\,\text{kg}\,\text{m}^{-3} \times 0.4\,\text{m}^2}\right)^{1/3} \sim (1000\,\text{m}^3\,\text{s}^{-3})^{1/3} = 10\,\text{m}\,\text{s}^{-1}.$$

I'd use the 0.5 m^2 instead of 400W... (1) 5s are "nicer" numbers for people to deal with in general [even though it doesn't matter at all] (2) you've already admitted to the distinct possibility that the power is an underestimate & now you want to make it smaller. (3) I'm more comfortable with the idea of the cyclist being a bit larger than with him/her using less power.

So the world record should be, if this analysis has any correct physics in it, around $10\,\mathrm{m\,s^{-1}}$ or 22 mph.

The world one-hour record – where the contestant cycles as far as possible in one hour – is 49.7 km or 30.9 mi. The estimate based on drag is reasonable!

I still have to do this conversion the long way (*3600/1600) every time. My intuition for speeds in m/s beyond a small scale just isn't very good yet.

I used to have that problem until I starting remembering a few different numbers in meters per second. for example, if you remember how fast people walk, run, a few common speed limits (30mph and 60 mph), and the speed of sound (those are the ones I know), you gain an intuition really easily by comparison

Although I think this is a powerful estimation tool and an interesting application, I'm not convinced of it's practicality. In athletics, the difference between 22 miles and 30.9 miles is astronomical.

that's way cool

So the world record should be, if this analysis has any correct physics in it, around $10\,\mathrm{m\,s^{-1}}$ or 22 mph.

The world one-hour record – where the contestant cycles as far as possible in one hour – is 49.7 km or 30.9 mi. The estimate based on drag is reasonable!

Plus you lowered the power guess, which is part of the lower estimate error.

I was going to say the same thing if we had left the power at 500W the answer would be even closer

Barely though. This only raises it to about 10.7 m/s, or 24 mph.

I still have problems with this... if you did all the same calculations with a runner instead of a biker, you're really calculating how fast the person is capable of sprinting exerting his/her maximum power. Assuming you get a value similar to how fast a person can actually sprint, that value does not come close to representing how far someone could run in one hour since humans can barely sprint over a minute a max speed.

I would argue though that most of a sprinters energy isn't lost to drag. Also, his power estimate was for an athlete's "endurance" power production.

Actually, drag is very important in sprinting. Thats why records have to be wind dated below a certain level to count. However, the running example probably wouldnt work because people tend to run on tracks where the wind would mostly cancel out. Therefore the biking example seems to make more sense.

The wind definitely doesn't cancel out (speaking from many years of track and field at a highly competitive level). The wind against you imposes a force, changing your form and slowing you down. The wind behind you, while it makes it somewhat easier to run, again changes your form and will thus tire you out. This is a similar effect to how running downhill for a good portion of a cross country course can be detrimental if one hasn't trained for it.

You're probably right about that. I would argue that the wind definitely does speed you up when its against your back, but probably not enough to offset the wind blowing against you on the opposite side of the track. As for running downhill, people generally do run much faster which is why there are restrictions on the net change in elevation in road races.

So the world record should be, if this analysis has any correct physics in it, around 10 m s^{-1} or 22 mph.

The world one-hour record – where the contestant cycles as far as possible in one hour – is 49.7 km or 30.9 mi. The estimate based on drag is reasonable!

It definitely does make you run faster in terms of instantaneous velocity, but you'll be more tired afterwords than you would be running a similar distance on level ground due to increased muscle involvement in slowing you down, as well as spending significantly more time with your feet on the ground (since you'll be striking more in front of your body than underneath your hips), thereby making your form less bio-mechanically efficient

Yes I agree that your muscles would be more sore from running inefficiently. But my point is that despite this your speed would be improved by the wind at your back or gravity pulling you down a hill. In that sense, it will "cancel" with the opposing force.

So if you aren't satisfied with the explanation of wind "canceling out" on a track, a better explanation would be that when running slower speeds the wind speed is less negligible (since v is in the equation for drag). That's why biking makes a better example.

But even Lance can't maintain 500W for more than 30 minutes, so there is some other source of error here. Is coasting helping more than we think, or is drag not as high as we think?

I don't really think drafting is why we messed up just because I think the record is a one person event. If what you meant by coasting is stopping the pedal motion, I think our calculation takes that into account because we're averaging power. Since stopping would slow you down and you'd lose ground to that air resistance, you'd have to pedal harder afterwards and your average power would be more or less the same.

That's an interesting question about sprinting, but the power delivered in sprinting is pretty different. When you cycle you really pushing the pedals down just as you are pushing down on the ground when you run up stairs. When you run forward, in some sense you are just falling forward, but I think there is some power lost due to that up and down motion.

Good to see we have a bunch of runners in the class. Thanks for the info guys! As for the question, I think the output power is high but the explanation is crisp and clean.

So the world record should be, if this analysis has any correct physics in it, around $10\,\mathrm{m\,s^{-1}}$ or 22 mph.

The world one-hour record – where the contestant cycles as far as possible in one hour – is 49.7 km or 30.9 mi. The estimate based on drag is reasonable!

side bar this!

this is still quite a bit larger than our estimate, especially considering we know some pretty strict bonds on what it could be, even without all the math

This is a great method for figuring this out. I definitely understand drag more. Can we apply this to viscosity.

Did you even talk about fleas? You said you would at the beginning of the section

Agreed. The fleas sounded intriguing. Maybe that was just a statement indicating how general this technique is? It certainly sounded like an example with fleas was coming.

Haha, I noticed that too. I was rather interested to see what you had to say.

It might be useful to succinctly restate a general way of calculating drag (to conclude the section)

I agree. For me, it would also be helpful to have a 1 or 2 line reminder of the estimations/simplifications we made to arrive at that final drag force equation.

Thats not a bad idea. I really like some of the closing paragraphs in previous sections. They really concluded the processes well. Even though the application of symmetry was pretty apparent, I don't think the word "symmetry" was ever used in this section.

I agree, a concluding paragraph would be nice. I can see that I will be using some of the information in this reading to complete the homework, but my thoughts on this section are a little scattered and unorganized. It would be nice to tie everything together at the end.

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

I have real experience with any of the physics behind flight because we don't cover it in 8.01...so this entire section is very cloudy for me.

I really like these examples in mechanical engineering, however, will we ever discuss any more comp sci/ee readings?

think this should have been mentioned earlier.

how would you modify this approach to calculate the minimum-energy speed for a car?

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

A general note: it would be really awesome to have all of the week's readings posted on Sunday, so we could sit down for an hour and do them all if that would be most convenient for us. It would allow us to gain a broader sense of the topic, since when we focus really hard on one at a time I, at least, tend to lose sight of the general principles being taught.

Also, another point: sometimes I've found myself reading the memos just to post things on them (to get it checked off), rather than reading them because I want to learn the material. I don't have a solution to this problem, but it seems to act counter to the principle of not rewarding people so they actually care about learning.

Read Section 3.5 for the memo on Thursday. One topic in lecture on Friday will be a home experiment (one for everyone) to refute a common but bogus explanation of lift.

Is there a way to undo an "I agree" or "I disagree" ?

yes, click on it again...it's a "toggle" field.

A recitation teacher in my freshman year 8.01 class described flight as air taking more time to travel over a wing (due to curvature) then to flow underneath it, creating lift.

I believe this is exactly what lecture will disprove on Friday.

i don't think lecture today "disproved" the theory. it just proposed a better, more testable one.

I don't know... it seems what we did in lecture is somewhat incompatible with that explanation.

I always wondered how far they can go without landing.

I did not see any analysis of how far birds can fly in this section. If none of the further sections mention this (I think they should, becuase bird flight is incredibly fascinating) I would recommend you remove the reference to how far can birds fly, if you don't answer it.

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

what are streamlines?

streamlines are kind of like electric field lines in E&M. they're lines that are always tanget to the fluid velocity.

what are streamlines?

Streamlines are the way that the air moves over and under the object, in this case, a bird. They help explain how the air moving past the bird both push the bird downward and upward.

I think this reads awkwardly. Maybe you could say "...and much more."

what are vortices? was this mentioned before in the class?

He has talked about vortices in class. it's the spinning flow of a fluid. it's when a fluid spins rapidly around a center point. think about stirring coffee or pouring water. sometimes you get little vortices in them.

missing "to", use conservation to estimate

I think it would also be helpful if we knew immediately how this relates to symmetry. Although the brief explanation of the example is great!

I really like estimation problems like these that are tied to physics. I see them naturally in physics problems a lot.

Is this actually how planes operate?

maybe tie back to real world... to minimize costs

maybe it's just me but i had to read this a few times to get it (i'm probably just slow), maybe having "minimum-energy" in quotes would help

to clarify, the following phrase confused me: minimum energy way. so i suggested that if it said "minimum energy" way, I'd understand what you were trying to say more immediately.

Invariant!

Yea! I never thought to think of flyijng as an invariant problem.

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

I didn't know this was true. Thats cool.

Well, 'constant altitude' means that vertical velocity is zero, so the vertical component of momentum can't possibly change.

Hm this makes sense. But when a plane is landing, why does the pilot curl the wings in (downward)–wouldnt that cause more deflection of air downwards?

I thought that was the point? By curling wings in and deflecting air downward, the plane gets some upward lift.

The deflection of flaps and increase in camber allows the plane to gain lift at a lower airspeed.

Because the air's velocity is lower, it's momentum is lower. So this allows the plane to compensate.

Yeah, I haven't be aware of this!

I had not realized until now that that is how lift actually works in a physics forces kind of view.

I don't think that this is how lift actually works. Maybe this is an approximation for an airfoil.

Actually, it is surprisingly accurate. Some cool pictures demonstrate this:

http://people.eku.edu/ritchisong/554images/downplane5.gif

http://www.aviation-history.com/theory/lift_files/fig6.jpg

Those are awesome images.

I found this explanation helpful as well.

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{\text{air}} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

This is much more satisfying of an explanation than Bernoulli's principle for why a plane stays in the air.

I don't really agree that it's a more satisfying explanation, but definitely a more creative explanation which I enjoy.

I don't think I could have come up with this explanation by myself.

we all know through physics that velocity in x is not related to y. Thus we had to come to the conclusion that although the turbines are thrusting us forward, something is keeping us from falling. That is Bernoulli.

Shouldn't this go after the 3 steps when we define the other terms?

I agree, I think that works. It doesn't matter a whole lot though... you still get the picture.

agreed... when I got to the next part, I had to look for this definition again

I also agree; the order seems a little off here.

Seems strange to introduce this in the first-person.

These steps aren't intuitive to me.

I think we can think about it like this: in order to fly we need to displace some air, but how much? We have to displace the air with a certain velocity to actually get airborne, and we want to know how much energy we need to use to do that.

So I thought this would be a very hard calculation. But I thought that it was explained very well. It especially helped me when you went through the dimensions at the end.

I think it's somewhat difficult to calculate the mass of air, because the volume is almost infinite

Why is the air deflected downward from a flat plate like the index card we used in lecture?

as the plane moves through the air, I believe it creates a vortex under and the air flows under the wing. The energy to displace the air should translate into energy which keeps the plane up.

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

this is also represented as "angle of attack"

Angle of attack is only part of this calculation, which has to have some other parameters like wing length in there.

As in how much air, or the angle of deflection?

Good point. I think this is referring to finding m.air, so it's what quantity.

obviously this is divide and conquer- where does symmetry tie in?

This is a very interesting way to think about this problem

Every step leads to the next one. I would have liked to see a tree here. Steps make it seem like we do 3 things when we really do 2 and combine them to find the third.

is this a function of the speed at which the airplane is traveling?

and also, is it a function of height?

I think that makes sense. I'm not too familiar with this type of a problem, but i imagine maintaining a particular angle of deflection will maintain the same lift, regardless of height.

It'd be cool to eventually see a tree for this- it would show yet again another tie between what we have done and what we are doing now

I totally agree. This reading has a lot of math and it would definitely be helpful seeing a tree to help tie it all in.

Divide and conquer hard at work!!

Why is kinetic energy imparted to the air matter? Doesn't the flight just depend on how much recoil there is?

Why is the area L^2 and not L*width of wing?

Why L^2 and not just the area of the wing's foot print?

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

Typo?

Why this area [is] L^2...

I understand the first L, and I understand s, but not the other L. If we looked at it like a box: L = width s = length L = Height? I don't understand the intuition

yeah i'm confused about where the L^2 is coming from too.

I think, based on the paragraph, that we are looking at how far the air is disturbed, rather than just where exactly the wings pass. As such, the air is disturbed in roughly a square, where on side of the square is the wings themselves (length L) and the other part is the region right behind where the plane just flew, also a length L back, leading the disturbed air area to be L^2

Oh, if that's what it's supposed to mean, that explanation made a lot more sense to me than what is in the actual text. Maybe explaining it in the text like that would be helpful?

I agree that that explanation made more sense. What I'm still confused about is why the width of the wing is not taken into consideration.

Seconded.... I was thinking about this for like 5 minutes and searching through all these comment boxes just to find a thread discussing this.

by this do you mean that as a plane is traveling there is basically a zone in front and behind of the plane where the surrounding air that is traveling with it acts as part of the wing?

This is very hard to follow; I feel like this should be expanded and made clearer.

Although the overall picture is different, this is the same area calculation that is used in using fluid dynamics to calculate lift force.

what's with all the random boxes. i hate this program...

i didn't think we were considering cross-sectional area at all. isn't the area under the wings we would look at instead?

Yeah, I think what he's saying in this sentence is that we are choosing this L^2 over the cross sectional area but stating that the reason for doing it this way is "subtle"

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

So the cross sectional area has no effect on air disturbance?

It seems so; this makes sense as the plane will be moving so the area over which the wing can deflect air is roughly a square.

Is this because vertical momentum is conserved?

I am not very clear about this, is it supposed to be an approximation?

Yes! That's why he used a tilde ( ) - it means "approximately"

This sentence confused me, as its a bit wordy and I'm not familiar with the physics of flight. Is this saying that the only important factor is the length of the wings, and not their width at all?

Yeah, this seems a little confusing–how does area not matter? why does it only depend on length? wouldn't a plane with wide wings deflect more air?

Maybe it only depends on where the most disruption is ocurring, which would be the leading edge of the wing. And thus the dominant contribution goes as L^2 instead of being proportional to the cross sectional area of the wing?

how do you know that? did you learn it from physics or general knowledge

Why? This doesn't really seem like an explanation. More just like another fact you happened to know.

I agree that this seems like a fairly obscure fact to know. I'd imagine it's derived from computing velocity*arbitrary unit of time, which probably gives a distance traveled of about L. However, regardless of how this was arrived at, it should be shown.

Good point. A beautiful picture of the air disturbance is given in the Wikipedia "Vortex" article. Equally good, the picture was made by a US government agency (NASA), so it has no copyright, which means I can use it in the textbook.

are we talking characteristic length for the equation here? then again, i don't know why its l^2 rather than 2l...

I'm not sure about the L, but I know it can't be 2L since we need an area (units of distance^2)

3.5 Flight

How far can birds and planes fly? The theory of flight is difficult and involves vortices, Bernoulli's principle, streamlines, and much else. This section offers an alternative approach: use conservation estimate the energy required to generate lift, then minimize the lift and drag contributions to the energy to find the minimum-energy way to make a trip.

3.5.1 Lift

Instead of wading into the swamp of vortices, study what does not change. In this case, the vertical component of the plane's momentum does not change while it cruises at constant altitude.

Because of momentum conservation, a plane must deflect air downward. If it did not, gravity would pull the plane into the ground. By deflecting air downwards – which generates lift – the plane gets a compensating, upward recoil. Finding the necessary recoil leads to finding the energy required to produce it.

Imagine a journey of distance s. I calculate the energy to produce lift in three steps:

1. How much air is deflected downward?
2. How fast must that mass be deflected downward in order to give the plane the needed recoil?
3. How much kinetic energy is imparted to that air?

The plane is moving forward at speed v, and it deflects air over an area L^2 where L is the wingspan. Why this area L^2, rather than the cross-sectional area, is subtle. The reason is that the wings disturb the flow over a distance comparable to their span (the longest length). So when the plane travels a distance s, it deflects a mass of air

$$m_{air} \sim \rho L^2 s.$$

The downward speed imparted to that mass must take away enough momentum to compensate for the downward momentum imparted by gravity. Traveling a distance s takes time s/v, in which time gravity imparts a downward momentum Mgs/v to the plane. Therefore

This has to be a vertical distance right? I guess the passage of the plan disturbs a rectangular prism of air that has dimensions s x L x L?

Also, I suppose most wings have a similar shape where the longitudinal length is some fraction (a) of the span, L, making the total area some fraction of the total area: aLˆ2.

Surely their area along with their span must play some role. Otherwise airplanes would have thinner wings, right? Someone who knows more please explain.

With rho being the density of air?

of the fluid through which the vehicle travels, which in this is case, is air.

Didn't you just say this above?

Yes, but I think that it is rephrased slightly, making a momentum to momentum comparison.

While it's clearly stated what 's' means in this context, when we start throwing around speeds and masses, it's pretty natural to expect 's' to mean seconds. I think using 'd' for distance would make this much easier to follow.

Not a big deal or anything, but you don't actually define M (though it's pretty obvious of course).

I'm slightly confused as to how you got this so an explanation would be nice. Instantaneous momentum is given by Mv, but Mgs/v seems to be F*t?

So he has Mgs/v –> meaning gs/v has to equal a velocity if you want this to be momentum. g is m/sˆ2 and s/v is s, giving you m/s – so it all works out in the end. Also, the units of momentum are N-s, so F*t also works (as it should)!

I actually really like this bit of unit analysis, thanks whoever posted it!

$$m_{air}v_{down} \sim \frac{Mgs}{v}$$

so

$$v_{down} \sim \frac{Mgs}{vm_{air}} \sim \frac{Mgs}{\rho vL^2 s} = \frac{Mg}{\rho vL^2}.$$

The distance s divides out, which is a good sign: The downward velocity of the air should not depend on an arbitrarily chosen distance!

The kinetic energy required to send that much air downwards is $m_{air}v_{down}^2$. That energy factors into $(m_{air}v_{down})v_{down}$, so

$$E_{lift} \sim \underbrace{m_{air}v_{down}}_{Mgs/v} v_{down} \sim \frac{Mgs}{v} \underbrace{\frac{Mg}{\rho vL^2}}_{v_{down}} = \frac{(Mg)^2}{\rho v^2 L^2} s.$$

Check the dimensions: The numerator is a squared force since Mg is a force, and the denominator is a force, so the expression is a force times the distance s. So the result is an energy.

Interestingly, the energy to produce lift decreases with increasing speed. Here is a scaling argument to make that result plausible. Imagine doubling the speed of the plane. The fast plane makes the journey in one-half the time of the original plane. Gravity has only one-half the time to pull the plane down, so the plane needs only one-half the recoil to stay aloft. Since the same mass of air is being deflected downward but with half the total recoil (momentum), the necessary downward velocity is a factor of 2 lower for the fast plane than for the slow plane. This factor of 2 in speed lowers the energy by a factor of 4, in accordance with the v^{-2} in E_{lift}.

3.5.2 Optimization including drag

The energy required to fly includes the energy to generate lift and to fight drag. I'll add the lift and drag energies, and choose the speed that minimizes the sum.

The energy to fight drag is the drag force times the distance. The drag force is usually written as

$$F_{drag} \sim \rho v^2 A,$$

where A is the cross-sectional area. The missing dimensionless constant is $c_d/2$:

Although I understand where this came from it may have been helpful to have a description beforehand, maybe even a visual, of how you planned to use the information you were looking for.

this is very clear

Why does the twiddle become an equal sign at the end of this?

It's nice to see this work out mathematically because the intuition at the beginning is that d should not affect the air velocity

Yeah I like this little proof here.

I agree, it's similar to when you calculate things like Force/Length but need to pick an arbitrary length in the beginning of the problem only to cancel it out later.

Why isn't this 1/2mv^2?

Maybe he doesn't think the 1/2 is significant for this approximation. I think in another section he left out a factor of 1/2 for a drag equation.

I think this has to do with the approximation. This came up in the last reading as well.

Other people hinted at this, but I think there should be a twiddle in front of this equation.

is this "s" seconds?

Never mind. I see we are now looking at total energy from the flight

It's distance

Though that brings up a good point - when I first read the comment I thought, "of course it's seconds! That's the SI unit." s as a distance variable is a little confusing when numbers are involved - maybe change it to d?

I agree. We used d for distance in the last memo. It's good to be consistent.

No, it's distance. Kind of confusing notation, right?

$$m_{air}v_{down} \sim \frac{Mgs}{v}$$

so

$$v_{down} \sim \frac{Mgs}{vm_{air}} \sim \frac{Mgs}{\rho v L^2 s} = \frac{Mg}{\rho v L^2}.$$

The distance s divides out, which is a good sign: The downward velocity of the air should not depend on an arbitrarily chosen distance!

The kinetic energy required to send that much air downwards is $m_{air}v_{down}^2$. That energy factors into $(m_{air}v_{down})v_{down}$, so

$$E_{lift} \sim \underbrace{m_{air}v_{down}}_{Mgs/v} v_{down} \sim \frac{Mgs}{v} \underbrace{\frac{Mg}{\rho v L^2}}_{v_{down}} = \frac{(Mg)^2}{\rho v^2 L^2} s.$$

Check the dimensions: The numerator is a squared force since Mg is a force, and the denominator is a force, so the expression is a force times the distance s. So the result is an energy.

Interestingly, the energy to produce lift decreases with increasing speed. Here is a scaling argument to make that result plausible. Imagine doubling the speed of the plane. The fast plane makes the journey in one-half the time of the original plane. Gravity has only one-half the time to pull the plane down, so the plane needs only one-half the recoil to stay aloft. Since the same mass of air is being deflected downward but with half the total recoil (momentum), the necessary downward velocity is a factor of 2 lower for the fast plane than for the slow plane. This factor of 2 in speed lowers the energy by a factor of 4, in accordance with the v^{-2} in E_{lift}.

3.5.2 Optimization including drag

The energy required to fly includes the energy to generate lift and to fight drag. I'll add the lift and drag energies, and choose the speed that minimizes the sum.

The energy to fight drag is the drag force times the distance. The drag force is usually written as

$$F_{drag} \sim \rho v^2 A,$$

where A is the cross-sectional area. The missing dimensionless constant is $c_d/2$:

I thought kinetic energy was 1/2 * mv^2, or are we just ignoring the 1/2 for reasons of approximation.

I think we are ignoring that factor of a half for approximation purposes.

I had the same thought, and he seems to do it consistently, so yes.

We're ignoring the 1/2 not because it's a small factor, but because it's a constant factor at all. The ' ' in the equation indicates we are just trying to find a proportional relationship between energy and the other variables so that we can infer the effects of altering one of those other variables (like v, m, etc.)

This is something I have come to use a lot when solving problems with various units. Its a very good sanity check.

I agree, especially in these types of problems. If you can see that the dimensions work out the problem ends up becoming simple.

I don't understand the wording of this sentence- are we saying that it's a squared force because Mg (a force) is squared or is it because Mg is a force that it must be squared?

I believe it's your first answer, since we're just trying to put some meaning to these terms.

I like that we checked the dimensions without having to revert all the way back to mass, length, and time. I tend to do that, and clearly getting to force * distance is the quick and easy way!

This is true, but not obvious at first glance. A quick note about dimensional analysis might help.

It's equivalent to the dimensions of the drag force we saw earlier, when you consider [L^2] = [A] (both meters^2). [rho*v^2*L^2]=[rho*v^2*A]=[Force]

I like this note.... dimensional analysis is always a helpful tool... especially in approximations.

s is easily confused with seconds. A final draft might want to use a more recognizable variable, like 'd' or 'x'

$$m_{air}v_{down} \sim \frac{Mgs}{v}$$

so

$$v_{down} \sim \frac{Mgs}{vm_{air}} \sim \frac{Mgs}{\rho v L^2 s} = \frac{Mg}{\rho v L^2}.$$

The distance s divides out, which is a good sign: The downward velocity of the air should not depend on an arbitrarily chosen distance!

The kinetic energy required to send that much air downwards is $m_{air}v_{down}^2$. That energy factors into $(m_{air}v_{down})v_{down}$, so

$$E_{lift} \sim \underbrace{m_{air}v_{down}}_{Mgs/v} v_{down} \sim \frac{Mgs}{v} \underbrace{\frac{Mg}{\rho v L^2}}_{v_{down}} = \frac{(Mg)^2}{\rho v^2 L^2} s.$$

Check the dimensions: The numerator is a squared force since Mg is a force, and the denominator is a force, so the expression is a force times the distance s. So the result is an energy.

Interestingly, the energy to produce lift decreases with increasing speed. Here is a scaling argument to make that result plausible. Imagine doubling the speed of the plane. The fast plane makes the journey in one-half the time of the original plane. Gravity has only one-half the time to pull the plane down, so the plane needs only one-half the recoil to stay aloft. Since the same mass of air is being deflected downward but with half the total recoil (momentum), the necessary downward velocity is a factor of 2 lower for the fast plane than for the slow plane. This factor of 2 in speed lowers the energy by a factor of 4, in accordance with the v^{-2} in E_{lift}.

3.5.2 Optimization including drag

The energy required to fly includes the energy to generate lift and to fight drag. I'll add the lift and drag energies, and choose the speed that minimizes the sum.

The energy to fight drag is the drag force times the distance. The drag force is usually written as

$$F_{drag} \sim \rho v^2 A,$$

where A is the cross-sectional area. The missing dimensionless constant is $c_d/2$:

what is the reasoning behind this? I thought it was the opposite

The equation shows v^2 in the denominator, meaning that - mathematically - its increase results in a decrease for E.

Is that interesting? Wouldn't this be logical?

I thought it was interesting and I don't know very much about planes, so I didn't think it was that intuitive.

I think its both logical (since it makes sense given how we think about planes) but it isn't intuitive.

So that's why you need to speed up in order to fly (and also the reason why I can't fly)

Birds don't fly that fast though so how do they generate all that lift? By flapping their wings?

their mass is smaller? Just a guess... I don't know much about flight/birds

their mass is smaller? Just a guess... I don't know much about flight/birds

is it right to think of this because youre flying faster so you're generating more lift per time step?

This seems to make sense however, as mentioned earlier it is hard for me to wrap my head around it. Gravity acts instantaneously so how can we say gravity has less time to pull the plane down. Is this one of those things where you have to blur your vision?

I had never thought of this before, it is an interesting way to think about the problem

Yea, I wouldn't have thought to think about the amount of lift needed this way. Intuitively it makes sense that a faster plane needs less lift, but actually calculating how much less is a different story.

is the E a function of time? I didn't think they were

$$m_{air} v_{down} \sim \frac{Mgs}{v}$$

so

$$v_{down} \sim \frac{Mgs}{v m_{air}} \sim \frac{Mgs}{\rho v L^2 s} = \frac{Mg}{\rho v L^2}.$$

The distance s divides out, which is a good sign: The downward velocity of the air should not depend on an arbitrarily chosen distance!

The kinetic energy required to send that much air downwards is $m_{air} v_{down}^2$. That energy factors into $(m_{air} v_{down}) v_{down}$, so

$$E_{lift} \sim \underbrace{m_{air} v_{down}}_{Mgs/v} v_{down} \sim \frac{Mgs}{v} \underbrace{\frac{Mg}{\rho v L^2}}_{v_{down}} = \frac{(Mg)^2}{\rho v^2 L^2} s.$$

Check the dimensions: The numerator is a squared force since Mg is a force, and the denominator is a force, so the expression is a force times the distance s. So the result is an energy.

Interestingly, the energy to produce lift decreases with increasing speed. Here is a scaling argument to make that result plausible. Imagine doubling the speed of the plane. The fast plane makes the journey in one-half the time of the original plane. Gravity has only one-half the time to pull the plane down, so the plane needs only one-half the recoil to stay aloft. Since the same mass of air is being deflected downward but with half the total recoil (momentum), the necessary downward velocity is a factor of 2 lower for the fast plane than for the slow plane. This factor of 2 in speed lowers the energy by a factor of 4, in accordance with the v^{-2} in E_{lift}.

3.5.2 Optimization including drag

The energy required to fly includes the energy to generate lift and to fight drag. I'll add the lift and drag energies, and choose the speed that minimizes the sum.

The energy to fight drag is the drag force times the distance. The drag force is usually written as

$$F_{drag} \sim \rho v^2 A,$$

where A is the cross-sectional area. The missing dimensionless constant is $c_d/2$:

I guess in theory this does make sense, but it's hard to wrap my head around this intuitively.

I agree. Seeing it written down on paper makes sense, but it's hard to picture it intuitively.

I disagree, I found the paragraph explanation very descriptive and exlpanatory. I think it is natural that this does not make intuitive sense, hence, the explanation.

I like this explanation. It is just like a projectile from 8.01 - except with a jet engine, but gravity still has the same effect.

This paragraph has been the best one so far. I think this was very informative and explained how lift is inverse to speed well.

That's really cool, I've never heard of that before. Maybe we could go into this a bit more- how fast would you have to go to not need any energy to produce lift??

The wording of this seems a little awkward. Maybe it would be better to go ahead and show the equations with the wording

a picture of all the forces acting on a plane would be nice here.

Agreed!

I agree but i think its important not to make any types of pictures too specific. I feel that in this class being able to extract the ideas from specific problems and apply them to any type of problem is crucial.

i thought "drag" would be in the "x" direction, while "lift" would be in the "y" direction, perpendicular to it. if this was the case. the magnitude of the vector sum of the two will not be the same as a scalar addition.

It seems like you can get around this issue by using "C" and the wingspan area instead of Cd and the cross sectional area.

Well, since the plane is moving in the 'x' direction, wouldn't lift have to have some effect on that as well? There is also airspeed velocity....what is generated by the engines

Drag and lift forces are vectors, but the energy used to overcome those forces is a scalar. So we can add Energy spent on lift and energy spent on drag to find total energy consumed.

that would be true if we were talking about summing forces, but not energy. Energy has no direction.

$$m_{air}v_{down} \sim \frac{Mgs}{v}$$

so

$$v_{down} \sim \frac{Mgs}{vm_{air}} \sim \frac{Mgs}{\rho v L^2 s} = \frac{Mg}{\rho v L^2}.$$

The distance s divides out, which is a good sign: The downward velocity of the air should not depend on an arbitrarily chosen distance!

The kinetic energy required to send that much air downwards is $m_{air}v_{down}^2$. That energy factors into $(m_{air}v_{down})v_{down}$, so

$$E_{lift} \sim \underbrace{m_{air}v_{down}}_{Mgs/v} v_{down} \sim \frac{Mgs}{v} \underbrace{\frac{Mg}{\rho v L^2}}_{v_{down}} = \frac{(Mg)^2}{\rho v^2 L^2} s.$$

Check the dimensions: The numerator is a squared force since Mg is a force, and the denominator is a force, so the expression is a force times the distance s. So the result is an energy.

Interestingly, the energy to produce lift decreases with increasing speed. Here is a scaling argument to make that result plausible. Imagine doubling the speed of the plane. The fast plane makes the journey in one-half the time of the original plane. Gravity has only one-half the time to pull the plane down, so the plane needs only one-half the recoil to stay aloft. Since the same mass of air is being deflected downward but with half the total recoil (momentum), the necessary downward velocity is a factor of 2 lower for the fast plane than for the slow plane. This factor of 2 in speed lowers the energy by a factor of 4, in accordance with the v^{-2} in E_{lift}.

3.5.2 Optimization including drag

The energy required to fly includes the energy to generate lift and to fight drag. I'll add the lift and drag energies, and choose the speed that minimizes the sum.

The energy to fight drag is the drag force times the distance. The drag force is usually written as

$$F_{drag} \sim \rho v^2 A,$$

where A is the cross-sectional area. The missing dimensionless constant is $c_d/2$:

Are we supposed to know all these formulas? Maybe you can make a file on the class site with all the formulas you've used so far in the book.

This was in the last chapter. I think he expects us to remember things from chapter to chapter

I think it would be nice to have a continuously updated list of different formulas that we have seen. Even one at the very end would be nice for reviewing for the final.

Why isn't this just written with the equation? Why is it mentioned afterward?

I agree. I'm not sure why you bothered to tell us that Fdrag is proportional to things and then immediately after tell us the exact formula for the force. It seems kind of redundant.

It's the same thing we do in class: figure out the relationships between the variables of a problem, then attempt to arrange them into a (semi-)formal equation.

yea same thing with the missing half in the kinetic energy equation. A usefull representation of how variables relate to each other is the more important then an exact answer.

$$F_{drag} = \frac{1}{2} c_d \rho v^2 A,$$

where c_d is the drag coefficient.

However, to simplify comparing the energies required for lift and drag, I instead write the drag force as

$$F_{drag} = C \rho v^2 L^2,$$

where C is a modified drag coefficient, where the drag is measured relative to the squared wingspan rather than to the cross-sectional area. For most flying objects, the squared wingspan is much larger than the cross-sectional area, so C is much smaller than c_d.

With that form for F_{drag}, the drag energy is

$$E_{drag} = C \rho v^2 L^2 s,$$

and the total energy to fly is

$$E \sim \underbrace{\frac{(Mg)^2}{\rho v^2 L^2} s}_{E_{lift}} + \underbrace{C \rho v^2 L^2 s}_{E_{drag}}.$$

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed $v_{optimum}$: the speed that minimizes the energy consumption for a fixed journey distance s. Going faster or slower than the optimum speed means consuming more energy. That extra consumption cannot always be avoided. A plane is designed so that its cruising speed is its minimum-energy speed. So at takeoff and landing, when its speed is much less than the minimum-energy speed, a plane requires a lot of power to stay aloft, which is one reason that the engines are so loud at takeoff and landing (another reason is probably that the engine noise reflects off the ground and back to the plane).

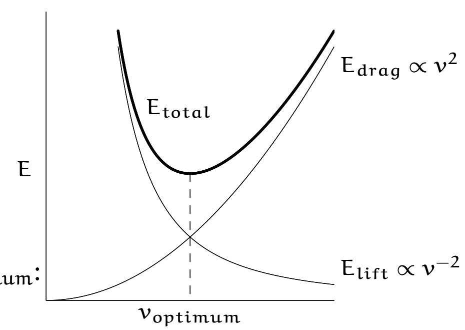

Where could we get something like this?

Why do we get hung up on including the constant coefficient C when we just nix the 1/2 from the KE equation so quickly?

According to wikipedia, a streamlined body has a drag coefficient of 0.04. This is a much more significant factor than a factor of 1/2 (that is, it changes the answer by a factor of ten).

Also, the drag coefficient mentioned above is c_d. Since C is much smaller even than that, C is yet more important.

One reason is that the 1/2 is the same for every object, whereas the C (drag coefficient) might vary a lot. And if want to understand the effect of shape on flight, then we need to keep the C around.

is the drag coefficient dependent on the cross-sectional area or L^2?

oh, hah, nevermind, answered in the next paragraph.

I'm still not sure what determines the drag coefficient and why it would need the additional coefficient of 1/2.

I think it has more to do with convention than need... I remember we discussed this in 2.006 but I forget where it came from.

Obviously the coefficient of drag could be expressed double the normal values in order to remove the 1/2, I'm pretty sure they didn't originally have the coefficient of drag and found that there was a factor of 1/2 different than just rho*v^2*A. Eventually they realized different geometries affect that coefficient, left the 1/2, and have been defining c_d in this way ever since. (A lot of geometries have a c_d of 1.)

$$F_{drag} = \frac{1}{2} c_d \rho v^2 A,$$

where c_d is the drag coefficient.

However, to simplify comparing the energies required for lift and drag, I instead write the drag force as

$$F_{drag} = C\rho v^2 L^2,$$

where C is a modified drag coefficient, where the drag is measured relative to the squared wingspan rather than to the cross-sectional area. For most flying objects, the squared wingspan is much larger than the cross-sectional area, so C is much smaller than c_d.

With that form for F_{drag}, the drag energy is

$$E_{drag} = C\rho v^2 L^2 s,$$

and the total energy to fly is

$$E \sim \underbrace{\frac{(Mg)^2}{\rho v^2 L^2} s}_{E_{lift}} + \underbrace{C\rho v^2 L^2 s}_{E_{drag}}.$$

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed $v_{optimum}$: the speed that minimizes the energy consumption for a fixed journey distance s. Going faster or slower than the optimum speed means consuming more energy. That extra consumption cannot always be avoided. A plane is designed so that its cruising speed is its minimum-energy speed. So at takeoff and landing, when its speed is much less than the minimum-energy speed, a plane requires a lot of power to stay aloft, which is one reason that the engines are so loud at takeoff and landing (another reason is probably that the engine noise reflects off the ground and back to the plane).

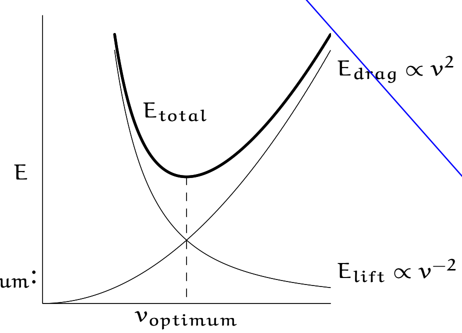

Does the capital C represent the 1/2*c_d? And if so, why do we consider the 1/2 here and not for kinetic energy?

The equation relating C and cd would be helpful to see here before C is substituted in. Is it C = (1/2)*cd*A/L^2 ?

Yes, I am confused why we introduced the 1/2 in the first equation but then have eliminated it here. It almost seems like the 1/2 was made up, and I don't know why.

Read the lines directly above in the reading. He explains it there.

is the drag that we are measuring the skin friction drag that occurs as air flows over the wing or the drag due to the air directly hitting and pushing against the front of the wing (momentum transfer). shouldn't both drags be considered?

I'm pretty sure the drag due to the air hitting the wing directly is quite low. The drag on the top of the wing has a larger effect on the bird because of the anatomy of birds and their ability to cut through the air well.

I don't understand the difference between using c_d and C

c_d is the coefficient in relation to the cross sectional area of an object, while C is a "modified" coefficient to use with the squared length of the wingspan. As mentioned, C is much smaller than c_d.

c_d is the coefficient in relation to the cross sectional area of an object, while C is a "modified" coefficient to use with the squared length of the wingspan. As mentioned, C is much smaller than c_d.

c_d=1/2*C

you are assuming that the drag generated by flow over the wing is much greater than the cross-sectional area to the point where the latter can be neglected?

Is this really just for most flying objects, or does it apply to all flying objects?

$$F_{drag} = \frac{1}{2} c_d \rho v^2 A,$$

where c_d is the drag coefficient.

However, to simplify comparing the energies required for lift and drag, I instead write the drag force as

$$F_{drag} = C\rho v^2 L^2,$$

where C is a modified drag coefficient, where the drag is measured relative to the squared wingspan rather than to the cross-sectional area. For most flying objects, the squared wingspan is much larger than the cross-sectional area, so C is much smaller than c_d.

With that form for F_{drag}, the drag energy is

$$E_{drag} = C\rho v^2 L^2 s,$$

and the total energy to fly is

$$E \sim \underbrace{\frac{(Mg)^2}{\rho v^2 L^2} s}_{E_{lift}} + \underbrace{C\rho v^2 L^2 s}_{E_{drag}}.$$

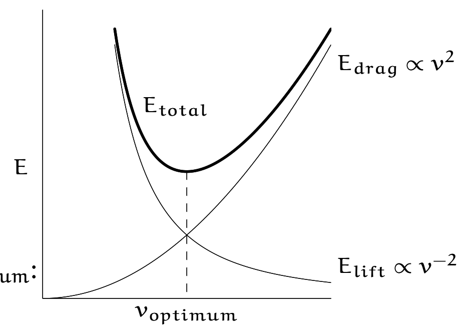

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed $v_{optimum}$: the speed that minimizes the energy consumption for a fixed journey distance s. Going faster or slower than the optimum speed means consuming more energy. That extra consumption cannot always be avoided. A plane is designed so that its cruising speed is its minimum-energy speed. So at takeoff and landing, when its speed is much less than the minimum-energy speed, a plane requires a lot of power to stay aloft, which is one reason that the engines are so loud at takeoff and landing (another reason is probably that the engine noise reflects off the ground and back to the plane).

I'm a little concerned about the usage of ' ' combined with the '+' sign in the right hand side of the equation. Unless you introduce a factor of C.Lift as a coefficient for the first term in the r.h.s., I think you're covering up the fact that here the constant C is also taking into account this lift coefficient.

The reason is that E CL*E.Lift + C*E.drag is not the same as E E.Lift + C*E.drag when you have already defined C.

I'm finding myself getting lost in all this math. It's not that I can't understand it, but it's so densely packed in here, that it's hard to concentrate. I'm not sure if you could get away with removing part of it, but it might help with flow and comprehension.

I think it's a necessary evil. You can't really explain things like lift and drag without resorting to mathematics.

We do need math, but the operations and substitutions could be presented more clearly or explained better. There are some quick leaps in here.

I too would like to see some clearly proofs, as well as an image of the plane. I do like the sequence of thinking however, moving from lift to drag and then combining the two.

I think a really well done diagram of the plane would be really helpful.

I too would like to see a diagram of the plane. Perhaps a force diagram showing what parts of the equation look like on the plane.

I agree, we have seen a lot of equations so far..and I'm kind of confused about which way (vertically or horizontally) the forces are acting on the plane.

A diagram would also be helpful in the beginning of the problem in terms of defining L and s.

Yes, I do think math is unavoidable at this here point now too. However, perhaps we could include a picture, that would make people feel that it is easier to understand, like when we talked about how pictures are more easily digested!

I agree I would definitely like to see a labeled force diagram of the plane here.

$$F_{drag} = \frac{1}{2} c_d \rho v^2 A,$$

where c_d is the drag coefficient.

However, to simplify comparing the energies required for lift and drag, I instead write the drag force as

$$F_{drag} = C\rho v^2 L^2,$$

where C is a modified drag coefficient, where the drag is measured relative to the squared wingspan rather than to the cross-sectional area. For most flying objects, the squared wingspan is much larger than the cross-sectional area, so C is much smaller than c_d.

With that form for F_{drag}, the drag energy is

$$E_{drag} = C\rho v^2 L^2 s,$$

and the total energy to fly is

$$E \sim \underbrace{\frac{(Mg)^2}{\rho v^2 L^2} s}_{E_{lift}} + \underbrace{C\rho v^2 L^2 s}_{E_{drag}}.$$

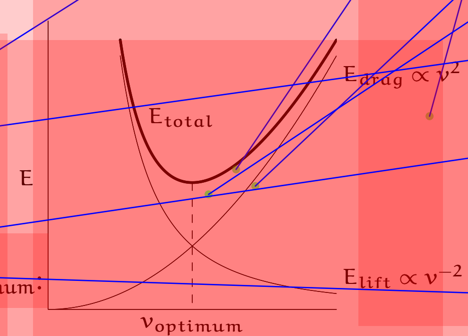

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed $v_{optimum}$: the speed that minimizes the energy consumption for a fixed journey distance s. Going faster or slower than the optimum speed means consuming more energy. That extra consumption cannot always be avoided. A plane is designed so that its cruising speed is its minimum-energy speed. So at takeoff and landing, when its speed is much less than the minimum-energy speed, a plane requires a lot of power to stay aloft, which is one reason that the engines are so loud at takeoff and landing (another reason is probably that the engine noise reflects off the ground and back to the plane).

so this means that the energy doesn't depend on the height at all? or is this the energy for a specific height

I think that this is ignoring the differences in the height, and therefore the density of the air, since it was previously estimated to be about 1 earlier. I would guess that this could make a big difference in the numbers. I think this would be a great point for the text to discuss.

I really like this graph- it really helps

So is this why birds fly in flocks, and why airliners are considering grouping up planes in fleets for trans continental and oceanic flights?

for birds, yes...I didn't know that they group planes for really long flights, but it does make sense.

I like this diagram. Because there are so many terms, this visualizes and simplifies the problem.

Good graph! This clarifies a lot, in terms of visualizing the shapes of the curves.

I agree very much as well. This graph really put things into perpective.

This is really cool to look at. I just always assumed the low speeds at takeoff and landing were mostly due to safety!

so are commercial planes designed so that this optimal cruising speed is just under mach 1?

I'm a little confused on how going under mach 1 effects the energy. What are you getting at?

Shouldn't we be calculating how fast birds actually fly, rather than the optimum speed?

$$F_{drag} = \frac{1}{2} c_d \rho v^2 A,$$

where c_d is the drag coefficient.

However, to simplify comparing the energies required for lift and drag, I instead write the drag force as

$$F_{drag} = C\rho v^2 L^2,$$

where C is a modified drag coefficient, where the drag is measured relative to the squared wingspan rather than to the cross-sectional area. For most flying objects, the squared wingspan is much larger than the cross-sectional area, so C is much smaller than c_d.

With that form for F_{drag}, the drag energy is

$$E_{drag} = C\rho v^2 L^2 s,$$

and the total energy to fly is

$$E \sim \underbrace{\frac{(Mg)^2}{\rho v^2 L^2} s}_{E_{lift}} + \underbrace{C\rho v^2 L^2 s}_{E_{drag}}.$$

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed $v_{optimum}$: the speed that minimizes the energy consumption for a fixed journey distance s. Going faster or slower than the optimum speed means consuming more energy. That extra consumption cannot always be avoided. A plane is designed so that its cruising speed is its minimum-energy speed. So at takeoff and landing, when its speed is much less than the minimum-energy speed, a plane requires a lot of power to stay aloft, which is one reason that the engines are so loud at takeoff and landing (another reason is probably that the engine noise reflects off the ground and back to the plane).

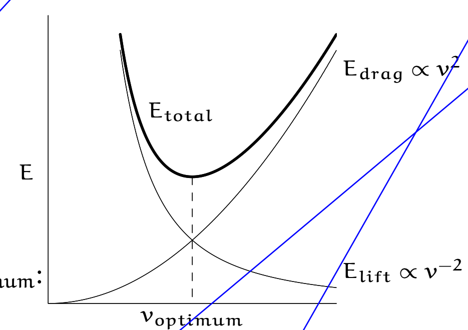

This is really interesting! I've always wondering what determined a plane's cruising speed.

Definitely! This course really does help answer some of the mysteries of life...

I always kind of assumed that planes fly close to as fast as they safely can.

The follow up question that immediately occurred to me...what determines their cruising altitude? Because the density changes as you move up or down...

I also always thought planes flew as fast as the safely could, it's interesting to see that they are maximizing their profit and could fly faster but it is not to their advantage.

Same for cars, like when there was a big gas crisis in the 70's people were told to stay at 55mph on the highway because that was the most efficient at the time.

55 is still the most efficient for most cars...i'm not sure about hybrids, but the rest are pretty well set there.

This is really interesting. Does this mean that different size planes travel at different speeds in order to minimize their energy consumption? I always thought that all airplane flights traveled at a roughly uniform speed. Does anyone have any insight or thoughts on this?

Yes! Similarly, different-sized birds have different cruising speeds. But for planes there is an additional constraint, which is that they follow one another on a flight path. So, it's not efficient to have planes with different cruising speeds. The 737, which otherwise would have had a different cruising speed than the 747, was redesigned so that it has the same cruising speed as a 747.

This is an interesting point, especially because it's not intuitive to me. I would think that going slower would consume less energy but now it makes sense why it would actually consume more.

you could also think about it in terms of a car, you use a lot more gas running around the neighborhood at 30-40 than you do racing down the highway at 55 mph...then again, that 30-40 uses about the same as 70-80.

nevermind, if we're doing planes.

$$F_{drag} = \frac{1}{2} c_d \rho v^2 A,$$

where c_d is the drag coefficient.

However, to simplify comparing the energies required for lift and drag, I instead write the drag force as

$$F_{drag} = C\rho v^2 L^2,$$

where C is a modified drag coefficient, where the drag is measured relative to the squared wingspan rather than to the cross-sectional area. For most flying objects, the squared wingspan is much larger than the cross-sectional area, so C is much smaller than c_d.

With that form for F_{drag}, the drag energy is

$$E_{drag} = C\rho v^2 L^2 s,$$

and the total energy to fly is

$$E \sim \underbrace{\frac{(Mg)^2}{\rho v^2 L^2} s}_{E_{lift}} + \underbrace{C\rho v^2 L^2 s}_{E_{drag}}.$$

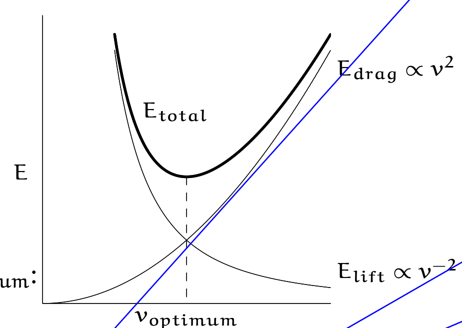

A sketch of the total energy versus velocity shows interesting features. At low speeds, lift is the dominant consumer because of its v^{-2} dependence. At high speeds, drag is the dominant consumer because of its v^2 dependence. In between these extremes is an optimum speed $v_{optimum}$: the speed that minimizes the energy consumption for a fixed journey distance s. Going faster or slower than the optimum speed means consuming more energy. That extra consumption cannot always be avoided. A plane is designed so that its cruising speed is its minimum-energy speed. So at takeoff and landing, when its speed is much less than the minimum-energy speed, a plane requires a lot of power to stay aloft, which is one reason that the engines are so loud at takeoff and landing (another reason is probably that the engine noise reflects off the ground and back to the plane).

Does this sort of logic to find v_opt apply to non-flying objects like cars too? I've heard it's more efficient to drive at some lower speed, but I was never sure what the forces were behind that.

I believe that's generally based on fuel efficiency of the engine. I've heard the most efficient spot is in top gear but with a low tachometer reading. It may have to be fine tuned to take air drag into account, though.

The drag term still holds (in some form) but the lift part doesn't, since a car doesn't need to sustain a certain velocity to stay up. It depends on the engine as mentioned, and lower rpm are more efficient, but the most efficient speed could theoretically be very slow, where drag was near zero. Then other operational definitions of efficiency might be needed, e.g., I need to get to work in less time than a day.

Is that always true? Is it true for supersonic speed?

I think that this is true for everything, except the planes designed to break supersonic speeds...they don't really care how much energy those take...

I don't think any plane actually 'cruises' at super sonic. It's like cruise control in a car I believe; and super sonic is like going 100 on the freeway.

Is this just saying that a plane is designed such that it is most efficient at its cruising speed...?

yes. If it wasn't then that would be very poor engineering. It's a good point that I never thought about.

Wouldn't takeoff be different from landing? Before landing, isn't the plane resisting less of gravity because it is heading towards the ground?

Interesting, that makes a lot of sense.

yeah really cool to learn that!

I never even thought about these two reasons. I always compared it to starting a car engine...but then the plane's engines are loud for such an extended period of time!

I never thought about this before, but it makes a lot of sense.

Aha! So at lower speeds, to compensate, the use of flaps increases the effective area of the wing. This shifts L^2 higher to compensate for a lower v^2.

The constraint, or assumption, that a plane travels at the minimum-energy speed simplifies the expression for the total energy. At the minimum-energy speed, the drag and lift energies are equal. So

$$\frac{(Mg)^2}{\rho v^2 L^2} s \sim C\rho v^2 L^2 s,$$

or

$$Mg \sim C^{1/2}\rho v^2 L^2.$$

This constraint simplifies the total energy. Instead of simplifying the sum, simplify just the drag, which neglects only a factor of 2 since drag and lift are roughly equal at the minimum-energy speed. So

$$E \sim E_{drag} \sim C\rho v^2 L^2 s \sim C^{1/2} Mgs.$$

This result depends in reasonable ways upon M, g, C, and s. First, lift overcomes gravity, and gravity produces the plane's weight Mg. So Mg should show up in the energy, and the energy should, and does, increase when Mg increases. Second, a streamlined plane should use less energy than a bluff, blocky plane, so the energy should, and does, increase as the modified drag coefficient C increases. Third, since the flight is at a constant speed, the energy should be, and is, proportional to the distance traveled s.

3.5.3 Explicit computations

To get an explicit range, estimate the fuel fraction β, the energy density $\mathcal{E}$, and the drag coefficient C. For the fuel fraction I'll guess $\beta \sim 0.4$. For $\mathcal{E}$, look at the nutrition label on the back of a pack of butter. Butter is almost all fat, and one serving of 11 g provides 100 Cal (those are 'big calories'). So its energy density is $9\,\text{kcal}\,\text{g}^{-1}$. In metric units, it is $4 \cdot 10^7\,\text{J}\,\text{kg}^{-1}$. Including a typical engine efficiency of one-fourth gives

$$\mathcal{E} \sim 10^7\,\text{J}\,\text{kg}^{-1}.$$

The modified drag coefficient needs converting from easily available data. According to Boeing, a 747 has a drag coefficient of $C' \approx 0.022$, where this coefficient is measured using the wing area:

$$F_{drag} = \frac{1}{2} C' A_{wing} \rho v^2.$$

nice, I like this assumption!

convenient.

The graph really helps in showing this is true.

Wait, isn't that a completely false statement? shouldn't their derivatives (with respect to v) be equal in magnitude and opposite in sign to minimize the sum of the energies?

Say, for example, that drag energy happened to be constant with velocity. Then there's no reason why you would want lift energy to equal drag energy, when you could just keep reducing it further.

Update: So I did the math, and coincidentally you actually get this exact same relation, but not for the reason stated in the text. Taking the derivatives of C1/v^2 and C2*v^2 to be equal and opposite you get:

2*C1/v^3 = 2*C2*v, which simplifies back to C1/v^2 = C2*v^2.

But this is only a coincidence because the velocity on the left side of the equation is raised to the opposite power as the velocity on the right side. I still believe that it is misleading to state that the energies will be equal without noting this fact.

I agree that this is a misleading way to state this. I wouldn't have been able to get there on my own with this explanation.

Yeah, I read this part of the text and did a total double take. He should really use the explanation provided by 12:20pm.

So once he explained it in lecture it made sense. Maybe that explanation should be made clear here....

Can someone reiterate that? I wasn't able to make lecture. Thanks.

This is a really good point. It makes it so much simpler and really understandable.

The constraint, or assumption, that a plane travels at the minimum-energy speed simplifies the expression for the total energy. At the minimum-energy speed, the drag and lift energies are equal. So

$$\frac{(Mg)^2}{\rho v^2 L^2} s \sim C\rho v^2 L^2 s,$$

or

$$Mg \sim C^{1/2} \rho v^2 L^2.$$

This constraint simplifies the total energy. Instead of simplifying the sum, simplify just the drag, which neglects only a factor of 2 since drag and lift are roughly equal at the minimum-energy speed. So

$$E \sim E_{drag} \sim C\rho v^2 L^2 s \sim C^{1/2} Mgs.$$

This result depends in reasonable ways upon M, g, C, and s. First, lift overcomes gravity, and gravity produces the plane's weight Mg. So Mg should show up in the energy, and the energy should, and does, increase when Mg increases. Second, a streamlined plane should use less energy than a bluff, blocky plane, so the energy should, and does, increase as the modified drag coefficient C increases. Third, since the flight is at a constant speed, the energy should be, and is, proportional to the distance traveled s.

3.5.3 Explicit computations

To get an explicit range, estimate the fuel fraction β, the energy density $\mathcal{E}$, and the drag coefficient C. For the fuel fraction I'll guess $\beta \sim 0.4$. For $\mathcal{E}$, look at the nutrition label on the back of a pack of butter. Butter is almost all fat, and one serving of 11 g provides 100 Cal (those are 'big calories'). So its energy density is $9\,\text{kcal}\,\text{g}^{-1}$. In metric units, it is $4 \cdot 10^7\,\text{J}\,\text{kg}^{-1}$. Including a typical engine efficiency of one-fourth gives

$$\mathcal{E} \sim 10^7\,\text{J}\,\text{kg}^{-1}.$$

The modified drag coefficient needs converting from easily available data. According to Boeing, a 747 has a drag coefficient of $C' \approx 0.022$, where this coefficient is measured using the wing area:

$$F_{drag} = \frac{1}{2} C' A_{wing} \rho v^2.$$

what? This seems nonsensical to me...missing words/typo?

"Simplifying the sum" = simplifying the sum that gives the total energy, written on the previous page.

Since E E-drag + E-lift, and these two are equal at min energy speed, E 2E-drag E-drag.

(I found this difficult to follow at first too.)

Why not just say we're looking at just 1/2 the total energy? Also, it would help to make explicit that we are substituting into the earlier equation for Edrag with this new relation to get Edrag in terms of Mg instead of p,v,L. "simplifying" seems like a fuzzy, relative term and it's used three times in quick succession here.

I also am pretty confused by what this sentence is trying to convey.

I think if you said "Instead of simplyfying the sum from the total energy equation you can simplify just the drag term..."

That may read better. I think the way that it is now is very confusing.

It seems strange to include C through all of this, given that it's unitless. Mg * s is an energy, correct? So would it be a twiddle here?

I like that we always analyze our results like this to ensure that they make intuitive sense.

what?

these are good sanity checks to verify that our equation makes sense.

I agree. In a lot of the problems, along with making sure dimensions match up, its important to understand that as the various variables change in value, the ultimate answer will change the way we expect.

WHy are we using these variables? where do they come in?

The constraint, or assumption, that a plane travels at the minimum-energy speed simplifies the expression for the total energy. At the minimum-energy speed, the drag and lift energies are equal. So

$$\frac{(Mg)^2}{\rho v^2 L^2} s \sim C\rho v^2 L^2 s,$$

or

$$Mg \sim C^{1/2}\rho v^2 L^2.$$

This constraint simplifies the total energy. Instead of simplifying the sum, simplify just the drag, which neglects only a factor of 2 since drag and lift are roughly equal at the minimum-energy speed. So

$$E \sim E_{drag} \sim C\rho v^2 L^2 s \sim C^{1/2} Mgs.$$

This result depends in reasonable ways upon M, g, C, and s. First, lift overcomes gravity, and gravity produces the plane's weight Mg. So Mg should show up in the energy, and the energy should, and does, increase when Mg increases. Second, a streamlined plane should use less energy than a bluff, blocky plane, so the energy should, and does, increase as the modified drag coefficient C increases. Third, since the flight is at a constant speed, the energy should be, and is, proportional to the distance traveled s.

3.5.3 Explicit computations

To get an explicit range, estimate the fuel fraction β, the energy density $\mathcal{E}$, and the drag coefficient C. For the fuel fraction I'll guess $\beta \sim 0.4$. For $\mathcal{E}$, look at the nutrition label on the back of a pack of butter. Butter is almost all fat, and one serving of 11 g provides 100 Cal (those are 'big calories'). So its energy density is $9\,\mathrm{kcal\,g^{-1}}$. In metric units, it is $4 \cdot 10^7\,\mathrm{J\,kg^{-1}}$. Including a typical engine efficiency of one-fourth gives

$$\mathcal{E} \sim 10^7\,\mathrm{J\,kg^{-1}}.$$

The modified drag coefficient needs converting from easily available data. According to Boeing, a 747 has a drag coefficient of $C' \approx 0.022$, where this coefficient is measured using the wing area:

$$F_{drag} = \frac{1}{2} C' A_{wing} \rho v^2.$$

By this, we still mean approximations, right?

Yes. He is now assigning "real" values to the symbols from before. Of course, the "real" values are just approximations like everything else.

This is one of my favorite parts of this course, we can see how approximations really do work in the real world where nature doesn't care about being "exact" It gives us a real place to start when we look at complexity instead of wasting out time trying to figure out every detail about a problem

Range of what?

yeah, a definition of what "explicit range" is would be nice here.

I think here by range he means the distance that can be traveled by a 747 without stopping, but I agree that a definition here would be nice.

What exactly does the fuel fraction describe in this case? I originally thought it was a measure of efficiency, but we treat engine efficiency separately later in the paragraph.

I'm also wondering this. What is it a fraction of?

I feel like beta should be explicitly defined. 'energy density' seems to me to be a common enough term, and C was previously defined.

Fuel fraction is the fraction of the plane's mass due to the fuel. This quantity coupled with the energy density and efficiency relates the amount of energy we can carry to the mass of the plane.

The issue is that there are two masses that are important – the mass of the plane affects the lift energy and the mass of the fuel affects the total energy available. The fact that these masses end up being on opposite sides of the fraction (so we only have to know their ratio, not their values) is not necessarily intuitive, so this abstraction to a fuel fraction is not obvious.

I'm confused as to what you need this estimates for, and where you got them.

interesting- but is that really that accurate? Shouldn't we use a car/gallons of oil instead?

hahha that's a good idea!

yea I agree it makes everything flow and make sense

The constraint, or assumption, that a plane travels at the minimum-energy speed simplifies the expression for the total energy. At the minimum-energy speed, the drag and lift energies are equal. So

$$\frac{(Mg)^2}{\rho v^2 L^2} s \sim C\rho v^2 L^2 s,$$

or

$$Mg \sim C^{1/2} \rho v^2 L^2.$$

This constraint simplifies the total energy. Instead of simplifying the sum, simplify just the drag, which neglects only a factor of 2 since drag and lift are roughly equal at the minimum-energy speed. So

$$E \sim E_{drag} \sim C\rho v^2 L^2 s \sim C^{1/2} Mgs.$$

This result depends in reasonable ways upon M, g, C, and s. First, lift overcomes gravity, and gravity produces the plane's weight Mg. So Mg should show up in the energy, and the energy should, and does, increase when Mg increases. Second, a streamlined plane should use less energy than a bluff, blocky plane, so the energy should, and does, increase as the modified drag coefficient C increases. Third, since the flight is at a constant speed, the energy should be, and is, proportional to the distance traveled s.

3.5.3 Explicit computations

To get an explicit range, estimate the fuel fraction β, the energy density $\mathcal{E}$, and the drag coefficient C. For the fuel fraction I'll guess $\beta \sim 0.4$. For $\mathcal{E}$, look at the nutrition label on the back of a pack of butter. Butter is almost all fat, and one serving of 11 g provides 100 Cal (those are 'big calories'). So its energy density is $9\,\text{kcal}\,\text{g}^{-1}$. In metric units, it is $4 \cdot 10^7\,\text{J}\,\text{kg}^{-1}$. Including a typical engine efficiency of one-fourth gives

$$\mathcal{E} \sim 10^7\,\text{J}\,\text{kg}^{-1}.$$

The modified drag coefficient needs converting from easily available data. According to Boeing, a 747 has a drag coefficient of $C' \approx 0.022$, where this coefficient is measured using the wing area:

$$F_{drag} = \frac{1}{2} C' A_{wing} \rho v^2.$$

I don't understand how butter can be related to jet fuel. The process where the fuel is converted to energy could be entirely different for foods than that of combustible liquids for all that I know.

I agree, I understand that fat is energy dense, but how do we know it is anywhere near as energy dense as jet fuel?

This seems like an abrupt transition, maybe explain the use/relevance of Beta and Epsilon....

I agree - I literally went 'woah!' as I read that first sentence. Where did all those variables suddenly come from?

I agree. Also it wasn't immediately clear to me that "an explicit range" meant we were going to calculate the maximum distance the plane could travel. Based on the previous section, I thought we were going to calculate a "range" of energy values used by the plane during flight.

Agreed! Where does beta and energy density play in? And typical engine efficiency? This went from making sense to making very little.

B*Eps = Energy/M or how much energy is our plane carrying per total mass B = fuel mass / M Eps = Energy from fuel / fuel mass All our energy is from fuel, so we can substitute B*Eps in for E/M in our formula for s.

The variables didn't bother me, as they were defined, but at the beginning I was also kindof confused what "explicit range" was exactly going to entail

The constraint, or assumption, that a plane travels at the minimum-energy speed simplifies the expression for the total energy. At the minimum-energy speed, the drag and lift energies are equal. So

$$\frac{(Mg)^2}{\rho v^2 L^2} s \sim C\rho v^2 L^2 s,$$

or

$$Mg \sim C^{1/2}\rho v^2 L^2.$$

This constraint simplifies the total energy. Instead of simplifying the sum, simplify just the drag, which neglects only a factor of 2 since drag and lift are roughly equal at the minimum-energy speed. So

$$E \sim E_{drag} \sim C\rho v^2 L^2 s \sim C^{1/2} Mgs.$$

This result depends in reasonable ways upon M, g, C, and s. First, lift overcomes gravity, and gravity produces the plane's weight Mg. So Mg should show up in the energy, and the energy should, and does, increase when Mg increases. Second, a streamlined plane should use less energy than a bluff, blocky plane, so the energy should, and does, increase as the modified drag coefficient C increases. Third, since the flight is at a constant speed, the energy should be, and is, proportional to the distance traveled s.

3.5.3 Explicit computations

To get an explicit range, estimate the fuel fraction β, the energy density $\mathcal{E}$, and the drag coefficient C. For the fuel fraction I'll guess $\beta \sim 0.4$. For $\mathcal{E}$, look at the nutrition label on the back of a pack of butter. Butter is almost all fat, and one serving of 11 g provides 100 Cal (those are 'big calories'). So its energy density is $9\,\text{kcal}\,\text{g}^{-1}$. In metric units, it is $4 \cdot 10^7\,\text{J}\,\text{kg}^{-1}$. Including a typical engine efficiency of one-fourth gives

$$\mathcal{E} \sim 10^7\,\text{J}\,\text{kg}^{-1}.$$

The modified drag coefficient needs converting from easily available data. According to Boeing, a 747 has a drag coefficient of $C' \approx 0.022$, where this coefficient is measured using the wing area:

$$F_{drag} = \frac{1}{2} C' A_{wing} \rho v^2.$$

I'm not sure I understand how exactly we go from the energy density of butter to the energy density of a plane?

I think the text is showing that we can estimate the energy density of jet fuel by calculating the energy density of a high-fat (high energy) substance like butter.

The Wiki article on Energy Density has a table that shows that Jet Fuel has almost the same energy density as biodiesel (vegetable oil) so I think the assumption the text makes is valid.

I think this needs much more justification in the text.

Right, it's energy density of any compact fuel source, and long-chain hydrocarbons (fatty acids in butter, alkanes in jet fuel), are pretty similar in energy density from C-C bonds. Butter does have about 20% water content (since it's an emulsion), so shortening or olive oil would also work.

I think it needs a transition sentence like the ones used above... something along the lines of "we can estimate the energy density of fuel based on that of butter"

Ok this explains it. Haha I was so lost. Questions like this require some research and can't be simply estimated. However, I think the way you attack this question is very interesting.

Why do the 'big calories' in butter help us in this calculation?

They help give a frame of reference for estimating the quantity of energy density.

This section seems kind of abrupt. Some introduction into what we are trying to calculate or accomplish here would be helpful.

Awkwardly phrased

"needs to be converted," or change from passive voice

The constraint, or assumption, that a plane travels at the minimum-energy speed simplifies the expression for the total energy. At the minimum-energy speed, the drag and lift energies are equal. So

$$\frac{(Mg)^2}{\rho v^2 L^2} s \sim C\rho v^2 L^2 s,$$

or

$$Mg \sim C^{1/2} \rho v^2 L^2.$$

This constraint simplifies the total energy. Instead of simplifying the sum, simplify just the drag, which neglects only a factor of 2 since drag and lift are roughly equal at the minimum-energy speed. So

$$E \sim E_{drag} \sim C\rho v^2 L^2 s \sim C^{1/2} Mgs.$$

This result depends in reasonable ways upon M, g, C, and s. First, lift overcomes gravity, and gravity produces the plane's weight Mg. So Mg should show up in the energy, and the energy should, and does, increase when Mg increases. Second, a streamlined plane should use less energy than a bluff, blocky plane, so the energy should, and does, increase as the modified drag coefficient C increases. Third, since the flight is at a constant speed, the energy should be, and is, proportional to the distance traveled s.

3.5.3 Explicit computations

To get an explicit range, estimate the fuel fraction β, the energy density $\mathcal{E}$, and the drag coefficient C. For the fuel fraction I'll guess $\beta \sim 0.4$. For $\mathcal{E}$, look at the nutrition label on the back of a pack of butter. Butter is almost all fat, and one serving of 11 g provides 100 Cal (those are 'big calories'). So its energy density is $9\,\text{kcal}\,\text{g}^{-1}$. In metric units, it is $4 \cdot 10^7\,\text{J}\,\text{kg}^{-1}$. Including a typical engine efficiency of one-fourth gives

$$\mathcal{E} \sim 10^7\,\text{J}\,\text{kg}^{-1}.$$

The modified drag coefficient needs converting from easily available data. According to Boeing, a 747 has a drag coefficient of $C' \approx 0.022$, where this coefficient is measured using the wing area:

$$F_{drag} = \frac{1}{2} C' A_{wing} \rho v^2.$$

how is it appropriate to look up this number, since this is an estimation class?

Probably. However, you can reasonably assume that something aerodynamic has a small drag coefficient, so an estimate would work too here (I would have guessed a drag coefficent of 0.01-0.1, which contains the actual value).

I understand that an aerodynamic object would have a small drag coefficient, but is there another way to arrive at this estimate if we don't have any intuition about what qualifies as a "small" coefficient? Or is it just something we should memorize?

I find it a little weird that we just looked up a number for the drag coefficient... but estimated the fuel's energy density via butter. I feel like if you're going to look up values like the drag coefficient, why not look up other numbers to get a more accurate answer anyway?

Sometimes you have to use easily available data.

It would be nice to mention that it's "measured using the wing area [in a different definition]:".

I agree that it's weird that we just have the coefficient of drag. Especially since we didn't use it for the problems in the previous section.

Alas, this formula is a third convention for drag coefficients, depending on whether the drag is referenced to the cross-sectional area A, wing area A_{wing}, or squared wingspan L^2.

It is easy to convert between the definitions. Just equate the standard definition

$$F_{drag} = \frac{1}{2}C' A_{wing}\rho v^2.$$

to our definition

$$F_{drag} = CL^2\rho v^2$$

to get

$$C = \frac{1}{2}\frac{A_{wing}}{L^2}C' = \frac{1}{2}\frac{l}{L}C',$$

since $A_{wing} = Ll$ where l is the wing width. For a 747, $l \sim 10\,\text{m}$ and $L \sim 60\,\text{m}$, so $C \sim 1/600$.

Combine the values to find the range:

$$s \sim \frac{\beta\mathcal{E}}{C^{1/2}g} \sim \frac{0.4\times 10^7\,\text{J}\,\text{kg}^{-1}}{(1/600)^{1/2}\times 10\,\text{m}\,\text{s}^{-2}} \sim 10^7\,\text{m} = 10^4\,\text{km}.$$

The maximum range of a 747-400 is 13,450 km, so the approximate analysis of the range is unreasonably accurate.

Comments on page 5

when would you know when to use which?

They're all interconvertible, just defined differently.

Yeah, in many of these problems there are various ways to go about getting the correct answer. Just like many physics problems, it all depends on how you define your directions.

I think there are too many random variables being introduced here. It makes it hard to follow/ keep track of what they all mean.

I agree, if I hadn't seen this material before I'd be very confused.

Why would you make wing width "l"? Why not a different variable like "w"?

but before we were using L^2... when do we use L^2 and when L*l?

Why did we know to multiply these?

I tried to address this on the previous page in a comment on Beta and Epsilon, but B*Eps = Energy/M in our equation rearranged to solve for s.

I am kinda lost here, how did you get this equation?

On just one stick of butter?! wow! (just kidding... but this is where having defined beta would add clarification to what this result means, I think.)

clarified butter or clarified beta?

I would be kind of interested to know what sort of distances most planes actually fly, as compared with their theoretical maximum range.

if im not mistaken, planes are required by law to carry atleast 20-25% extra fuel in case of emergencies and/or the need to divert to another airport. this is a good way to approximated the distances most planes would fly at based on their max range, as the range of planes various by their design points.

Alas, this formula is a third convention for drag coefficients, depending on whether the drag is referenced to the cross-sectional area A, wing area A_{wing}, or squared wingspan L^2.

It is easy to convert between the definitions. Just equate the standard definition

$$F_{drag} = \frac{1}{2}C' A_{wing}\rho v^2.$$

to our definition

$$F_{drag} = CL^2\rho v^2$$

to get

$$C = \frac{1}{2}\frac{A_{wing}}{L^2}C' = \frac{1}{2}\frac{l}{L}C',$$

since $A_{wing} = Ll$ where l is the wing width. For a 747, $l \sim 10\,\text{m}$ and $L \sim 60\,\text{m}$, so $C \sim 1/600$.

Combine the values to find the range:

$$s \sim \frac{\beta\mathcal{E}}{C^{1/2}g} \sim \frac{0.4\times 10^7\,\text{J}\,\text{kg}^{-1}}{(1/600)^{1/2}\times 10\,\text{m}\,\text{s}^{-2}} \sim 10^7\,\text{m} = 10^4\,\text{km}.$$

The maximum range of a 747-400 is 13,450 km, so the approximate analysis of the range is unreasonably accurate.

Maybe I've forgotten after reading 5 pages, but I don't remember at the beginning anything about trying to find the max range of the plane. The whole section seemed more like an exercise in math than trying to make a specific point or reach a certain conclusion. Maybe in the opening paragraph include a sentence about what specific values you want to calculate- which necessitates the calculations of lift and drag.

That's unfortunate. It took some paying attention, but I thought the section was pretty instructive.

I also think this should be thrown in at the beginning as a "goal" so we know why we care about lift/drag in the first place. I think it would give the chapter more of a direction.

Well, the very first sentence is: "how far can birds and planes fly?", which gives the whole section a goal. It should perhaps be emphasized a bit better, since it's so easily overlooked.

I just thought there were too many equations throughout the way that got me a little lost.

I agree, it's easy to get lost in all of the equations. It would be helpful if we could have a short recap of the key equations and processes just so we can put everything together.

I'll admit that I lost interest in the particulars of these calculations. What does this all have to do with symmetry again? I think it'd be a perfect place to draw things back to the topic of the chapter.

I agree, reading this section seemed like just reading some interesting information with lots of equations. It would be nice to know why I should pay attention to the details in all the equations given.

I think I'm a little lost after a once through read on this one; Somehow, the beginning and the end didn't work for me; it hasn't come together.

Can we more explicitly relate this back to the energies to generate lift and fight drag? That would tie this into the previous sections better.

How can it be unreasonably accurate? I guess all the estimation errors just happened to cancel out

why is it 'unreasonably' accurate?

why is it 'unreasonably' accurate?

Alas, this formula is a third convention for drag coefficients, depending on whether the drag is referenced to the cross-sectional area A, wing area A_{wing}, or squared wingspan L^2.

It is easy to convert between the definitions. Just equate the standard definition

$$F_{drag} = \frac{1}{2}C' A_{wing}\rho v^2.$$

to our definition

$$F_{drag} = CL^2\rho v^2$$

to get

$$C = \frac{1}{2}\frac{A_{wing}}{L^2}C' = \frac{1}{2}\frac{l}{L}C',$$

since $A_{wing} = Ll$ where l is the wing width. For a 747, $l \sim 10\,\text{m}$ and $L \sim 60\,\text{m}$, so $C \sim 1/600$.

Combine the values to find the range:

$$s \sim \frac{\beta\mathcal{E}}{C^{1/2}g} \sim \frac{0.4 \times 10^7\,\text{J}\,\text{kg}^{-1}}{(1/600)^{1/2} \times 10\,\text{m}\,\text{s}^{-2}} \sim 10^7\,\text{m} = 10^4\,\text{km}.$$

The maximum range of a 747-400 is 13,450 km, so the approximate analysis of the range is unreasonably accurate.

I don't really like the word 'unreasonably' at the end. To me, it carries the idea that we shouldn't be this close, but maybe we're just that good at making approximations (or more likely, our errors canceled).

Yeah I was sort of confused about what made it unreasonable...

Yes, aren't you expecting to be accurate? If you were not accurate, you probably would not have put it in the book.

was it a typo? Maybe use the word reasonably?

I'm still a little iffy on this as well, but maybe he just means that we are much closer than would be reasonably expected. I am often impressed with how the errors in this book cancel out to our advantage. Is this from careful choice of approximations or just plain luck?

why is it 'unreasonably' accurate?

I'm not quite sure what lesson this section was supposed to teach. I was able to follow and see how to calculate the energies needed for flight and gain a better understanding of the phenomena, however, I am not quite sure what the take home message is here?

4
Proportional reasoning

4.1 Flight range versus size 68
4.2 Mountain heights 70
4.3 Animal jump heights 72
4.4 Drag 76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

I think it'd be nice if you showed some dimensional analysis here

Despite being so short, this was a fairly difficult (to stomach without objecting to everything) section. I guess, we can just be glad that this class is set up such that all of our disbelief in methods and confusion can be brought to class the next day, where it can be answered in person.

why is it the square root of C?

I completely agree I don't really understand how this makes sense. Also in class we included a fact for efficiency...doesn't that matter anymore? Also are you say 's' is the approximately 1 across the plane industry or for each plane

This section was much easier to read than the last. It was very concise and that really helped me to follow it.

did this animal have a simpler estimate of fat as our example? Also wouldn't the weight/fat distribution matter?

I found this section a lot easier to read and understand than previous sections.

4
Proportional reasoning

4.1	Flight range versus size	68
4.2	Mountain heights	70
4.3	Animal jump heights	72
4.4	Drag	76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

Read the introduction and section 4.1 (more on flight) for Sunday's reading memo.

I think you need to use the phrase proportional reasoning more. It still feels like we're in older chapters.

what philosophy do you refer to in this statement? It's ambiguous.

I"m not really sure what this phrase means

I also think that there's so kind of grammatical mistake here...

I like the build-up, but I agree with the above two, the end of the sentence doesn't make any sense. What are you referring to? What philosophy...symmetry?

I took it to mean "one implementation of symmetry", but I'm not sure if we've seen a different implementation of symmetry that doesn't wring out complexity.

Oh thanks for the comment. I was confused until I took philosophy to mean symmetry.

I feel it becomes clearer when distance is proportional to time is written out shortly after.

Yeah, it's not that clear but I think he means we just use symmetry once, whereas for other things, like divide and conquer, or abstraction, we might have to use them multiple times in order to actually make things clearer

in = is?

No, it should be "in". Take out the qualifiers: Symmetry wrings out complexity and reasoning IN one implementation...

no it should be is. this is a section about proportional reasoning, so therefore one would not want to "wring" it out.

4
Proportional reasoning

4.1 Flight range versus size 68
4.2 Mountain heights 70
4.3 Animal jump heights 72
4.4 Drag 76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

i really hate this software. if you click anywhere it goes crazy and all these boxes pop up, it doesn't scroll like a normal pdf, and makes you do a hell of a lot of extra work.

> More details would be very helpful. There were similar comments a few readings back, and Sacha worked on it. Clearly more needs to be done, and specific descriptions of what goes wrong would help that along.
>
> For example, what scrolling doesn't work? (I use the scrollbar on the right and that seems to work.)

I think it's much improved and works fine - the main gripe i had earlier was that when you were typing and had a typo, you couldn't really click in your text box to fix it. But now that's been resolved.

"without accelerating" might be more intuitive, here.

I had to re-read this sentence to understand it, i think its talking too generally.

This paragraph seems very unclearly worded.

Could you be more specific like the example you listed right before?

This is clever, since this is basically like previously choosing an invariant. Here, we want to find two factors that are relative or proportional for reasons of comparison, simplifying two into one.

maybe "distance to time ratio"

"contributes to" might be better here.

Aren't these called extensive something?

4
Proportional reasoning

4.1 Flight range versus size 68
4.2 Mountain heights 70
4.3 Animal jump heights 72
4.4 Drag 76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

What if this were re-emphasized: "However, speed, the ratio of distance to time, is invariant". Or something similar?

I like this a lot better, mentioning the invariant. Given the previous sections it makes more sense

Definitely - with no forces on it, although meaning the same thing, is somewhat harder to understand.

I like the way this would tie in the previous section since we just went through learning invariants

The beginning of this section talks about symmetry too much. This is a whole new topic of proportional reasoning, but it feels like we're still on symmetry. Judging by everyone's comments thus far, I think it confused everyone else too because all the comments are asking what this section has to do with symmetry.

Perhaps an overall description of the topic would be nice, before jumping into specifics.

Yeah, after finishing this intro I have no idea where the next section is going.

what this section has to do with symmetry: symmetry reduces complexity just like proportional reasoning

I think a lot of this confusion would be resolved by explaining the first sentence more.

I think you're trying to present something which in simple in a complex manner. Speed is useful because it relates to two other useful things, distance and time.

independent of what time? this seems contradictory.

he means its invariant in this problem–with constant speed, the velocity (the derivative) does not have a "t" term; it's constant, so it doesnt depend on time

That's a better explanation of what is meant by this- the reading as it is now is confusing because it sounds like it's saying that distance/time is independent of time which doesn't make sense

I definitely found this confusing. How can he say that velocity=distance/time, and in the same sentence say velocity is independent of time.

4

Proportional reasoning

4.1	Flight range versus size	68
4.2	Mountain heights	70
4.3	Animal jump heights	72
4.4	Drag	76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

This seems contradictory with the beginning of the sentence. I understand what you're saying, but I still had to stop and think about it for awhile.

It confused me too..I think if the ratio distance/time wasnt there then I would have got it.

Yeah I agree - the ratio shows a clear dependence on time, and then you say there is no dependence on time (which is true) but it does throw something off.

Doesn't your proportionality statement also conflict with this?

I don't think it necessarily contradicts. We're told that there's a constant relationship between distance and time therefore speed is constant. That doesn't really contradict the proportionality statement because distance will be proportional to time according to the speed constant.

@12:09 : it's not that it actually contradicts, just that it takes rereading it a few times to realized that it's not.

It just didn't make sense at all to me...

I don't quite see why this makes it simpler than either distance or time.

On some level, this is rather silly. Time is very simple, at least at the level the average reader would take it. Perhaps a rephrasing emphasizing that we make the abstraction of velocity to simplify the relation between distance and time?

I don't really get this because it seems like distance would be proportional to 1/time or vice versa?

distance = velocity * time, so at constant velocity the distance is proportional to the time.

it doesn't matter what fraction it is, but only that if one changes, the other makes a corresponding change. that's proportionality.

the symbol looks more lik a symmetry sign than a proportional sign

I thought it was standard notation for proportional?

do we naturally use this symmetry when we say that something is x hours away, as opposed to a distance?

4
Proportional reasoning

4.1 Flight range versus size 68
4.2 Mountain heights 70
4.3 Animal jump heights 72
4.4 Drag 76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

This is an interesting way to look at velocity that does a good job of showing how we use symmetry in our everyday lives

Shouldn't this say "change for how long the object moves"?

this doesn't exactly jump out as me as symmetry.

yeah I'm a little confused about what the symmetry is

should be "moves"

Should this be: how long the object moveS?

what do you mean? How does this not change the conclusions?

I think the idea is that you are making a quick approximation and regardless of this, you will be left with a conclusion that can accurately answer your question. Not every answer we come upon will be completely accurate, but at least it will be accurate enough to make a solid conclusion.

This entire paragraph seemed to confuse what I thought was a very simple concept... the ratio of distance/time as an invariant. I'm still not sure how this is a symmetry operation.

I agree - I feel like it is trying to do with words what is intuitive for many of us. Maybe show how exponents are related on both sides of the equality? Above, distance is proportional to time. Maybe it would be easier to explain that distance is the only thing that changes because they are directly related to one another with equal powers?

There has to be a better way to word this entire section. I had to reread it twice to get the overall meaning. The idea is very simple: "Applying dimensional analysis to invariants produces useful result. Eg, knowing that speed is invariant means that you know the distance someone has traveled in a given time."

I agree that this paragraph is really abstract and complex. Thanks for the summary.

4
Proportional reasoning

4.1 Flight range versus size 68
4.2 Mountain heights 70
4.3 Animal jump heights 72
4.4 Drag 76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

I feel like this sentence was written very quickly - everything before the comma and everything after the comma basically say the same thing using similar terminology.

I think this whole section might have been written speedily, there are more grammatical mistakes here than in many other sections we've read

haha is that supposed to be a joke?

I just have to say. After reading the first 14 memos, I am very impressed with the amount of varied knowledge you have. The only problem is that readers probably won't have the same amount, which makes reading this difficult. Also, I'm very glad that the entire class reads and comments. It helps me understand the examples which in turn help me understand the concepts you are trying to teach.

What is the mean-square distance?

I think this is just the rms of the distance

It's the average value of the square of the distance traveled over many such random walks.

It's equivalent to the variance of distance traveled for a random walk where the expected distance traveled is 0 (keeping in mind that the distance could be negative or positive, depending on direction, in a 2D example).

why walks and diffusion? These seems like pretty random examples to use. Is there anything more general that we could use? Or one good example to follow through?

They're both examples of gaussian (or normal) distributions. The thing about this type of distribution is that pretty much anything is gaussian.

Yeah normal distributions are extremely useful. Although not always accurate, they are great for approximations.

I believe this is a misspelling.

"travelled" is more common in British English, and "traveled" is more common in American English. http://www.future-perfect.co.uk/grammartips/grammar-tip-travelled-traveled.asp

4

Proportional reasoning

4.1 Flight range versus size 68
4.2 Mountain heights 70
4.3 Animal jump heights 72
4.4 Drag 76

Symmetry wrings out excess, irrelevant complexity, and proportional reasoning in one implementation of that philosophy. If an object moves with no forces on it (or if you walk steadily), then moving for twice as long means doubling the distance traveled. Having two changing quantities contributes complexity. However, the ratio distance/time, also known as the speed, is independent of the time. It is therefore simpler than distance or time. This conclusion is perhaps the simplest example of proportional reasoning, where the proportional statement is

$$\text{distance} \propto \text{time}.$$

Using symmetry has mitigated complexity. Here the symmetry operation is 'change for how long the object move (or how long you walk)'. This operation should not change conclusions of an analysis. So, do the analysis using quantities that themselves are unchanged by this symmetry operation. One such quantity is the speed, which is why speed is such a useful quantity.

Similarly, in random walks and diffusion problems, the mean-square distance traveled is proportional to the time travelled:

$$\langle x^2 \rangle \propto t.$$

So the interesting quantity is one that does not change when t changes:

Each individual sentence make sense, but I think I'm still missing the big picture.

So far, he's been talking about how in these distance/time/velocity problems, there are invariants–for the "no external forces" case, it's velocity that's invariant, and for the "random walks" and "diffusion" cases, the mean-square distance travelled is always proportional to t (meaning that the invariant is x^2/t

Ah...thank you for that.

wouldn't abs(x) work too?

Perhaps? Is proportion in this case only linear, or are you saying only that one quantity depends on the other.

how is this related to symmetry...? seems more like using an invariant...

i agree. i dont really get all these categories and divisions and things.

what do the brackets mean? the bar means average, right?

I'm confused as well. But I think the important idea is that it doesn't always have to be linear.

This is probably a dumb question, but is it squared so going backwards is still considered as distance traveled or is there another reason?

Exactly. We square it to get something like total distance traveled and not displacement. Ordinarily I would expect to see root mean-squared and not just mean-squared though.

"Interesting quantity" means "invariant" here, yes? It would be better just to say so.

How did we come to define this as an invariant from the proportion?

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{\text{tank}}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{\text{tank}} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{\text{plane}}}{s_{\text{bird}}} \sim \frac{\beta_{\text{plane}}}{\beta_{\text{bird}}} \frac{\mathcal{E}_{\text{plane}}}{\mathcal{E}_{\text{bird}}} \left(\frac{C_{\text{plane}}}{C_{\text{bird}}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

Comments on page 2

Can you define what exactly is this "interesting quantity"? I understand what you're trying to say, but it's a little strange to read.

there is no definition for this "interesting quantity" other than the fact that it's the invariant. i guess he calls it "interesting" because there's no real name for it–it's just the mean square distance travelled divided by time (since they're proportional, this is always contant)

I am not understanding this quantity of not changing with time. Can you elaborate more in lecture?

I think I'm pretty confused about this entire section, but what is this inparticular meant to mean...?

I don't really understand this section either. How does the "interesting quantity" not change?

because the diffusion constant is an intrinsic property of a fluid, it doesn't vary with time or distance

I think explaining the concept with an example problem might make this material a little less confusing.

This 'interesting' quantity is not made clear; I feel like it's just being added and this part of the book is just vague and confusing.

I never thought about the diffusion constant in this context as an "interesting quantity," but now it makes sense why its important.

by that definition all constants are important, not just this one.

So it varies with each type of material?

Yes, I think that diffusion definitely varies with the type of material.

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{tank}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{tank} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{plane}}{s_{bird}} \sim \frac{\beta_{plane}}{\beta_{bird}} \frac{\mathcal{E}_{plane}}{\mathcal{E}_{bird}} \left(\frac{C_{plane}}{C_{bird}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

It might be nice to add an additional example or something. I feel like the section is easy to forget since the next section is so interesting.

I think it's ok the way it is...I read this section as an introduction to the chapter, so it doesnt matter if stuff gets a little forgotten here since the other sections will drive the point in.

This example seemed like it was gonna be explored further, but it just ended here. Perhaps it could have been incorporated into the actual sections.

What is the diffusion constant supposed to mean/represent?

I would like a little more information on this constant and maybe a couple common values to see how they compare. I think it would be helpful and interesting even if it doesn't completely add to proportionality, and continue in the theme of everyday knowledge that this book follows

Not sure where this is going. Need a simpler intro and focus on examples to prove your point. All your other examples are very good.

Maybe some introduction to this would be nice. Like "now we will look at how to to apply this to the range of a plane on a full tank of gas.."

What do you mean by range? How far a plane can go when its tank is full? On a gallon of gas? Requiring the least energy? Its somewhat ambiguous.

I believe that when he says range, he is referring to the maximum distance a plane can travel on a full tank of fuel.

I agree, but I think it would be useful to put that in there. My guess is it is "one plane, full capacity, full tank of gas, how far can it go"

But even there there's an important question of load - an unloaded large plane may travel farther than an unloaded small plane given the larger plane's larger fuel tank. However, that larger plane carries more weight when loaded, which would presumably effect the range.

i understand that this makes the approximation a lot easier, but it doesn't seem like it's true that a fighter plane has the same shape as a passenger plane like the 747, and I feel like this will have a huge impact on the resulting conclusions

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{\text{tank}}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{\text{tank}} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{\text{plane}}}{s_{\text{bird}}} \sim \frac{\beta_{\text{plane}}}{\beta_{\text{bird}}} \frac{\mathcal{E}_{\text{plane}}}{\mathcal{E}_{\text{bird}}} \left(\frac{C_{\text{plane}}}{C_{\text{bird}}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

Didn't know this equation.

Actually, I think you might.

If you notice the " ", it suggests that rather than being an explicit equation, it's just a statement of proportionality.

Here, it is just saying that the Energy of a plane to fly a distance is merely related (somehow or another) to a constant C, its Mass, the gravitational constant, and the distance traveled.

This seems very reasonable and I would bet that if you were asked what factors does the energy required to fly a certain distance depend on, you would come up with these factors. After all, a heavier object will require more energy. Traveling longer will also require more energy. And likewise, stronger gravity will require more energy to overcome it.

In this way, you don't have to think of this expression as an "equation" of the intimidating sort. Just think of it as an approximate relation that summarizes the various factors which influence energy.

It might help to explain where this comes from...just have a sentence or so.

I think this was covered in the previous section

It would be nice if some of the variables were explicitly stated - you mention C is the drag coefficient later, but it probably should be mentioned here.

I agree. Even if they've been mentioned in previous readings, it would be nice if we could be reminded again.

It might be nice to have C defined. (just realized it was defined lower down, maybe just define it when you introduce it?)

I think it was defined in an earlier section, if you assume continuity in the book across its sections.

Even with that assumption C can stand for more than one thing in many science/engineering examples. While the drag coefficient makes the most sense, I'd like to see it defined again.

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{tank}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{tank} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{plane}}{s_{bird}} \sim \frac{\beta_{plane}}{\beta_{bird}} \frac{\mathcal{E}_{plane}}{\mathcal{E}_{bird}} \left(\frac{C_{plane}}{C_{bird}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

what do these variables mean?

I am also confused. But I think 'M' is mass and 'g' is the gravitational constant?

yeah, this equation als gave me pause, maybe mention the coefficients, like C, closer to the beginning perhaps?

maybe go over what each coefficients mean? It is not completely obvious from the equation

What is s in this equation?

Look at the two words before the section you highlighted.

s is distance in this equation

does s refer to distance here? or seconds? It seems like a random letter to choose.

distance. distance is often denoted 's'.

Perhaps, but I am also used to distance as "d" and "s" as seconds...

I think you're confusing the variable with its units. 'd' and 's' are variables that have units of distance (meters, for example). Time is often 't' or 'T', and that has units of seconds, or 's'. 's' as a variable almost never represents time.

It does state that "s" is distance right above, but I agree that it tends to get confusing especially when the "s" is at the end.

He's used 's' to denote distance in previous sections.

is this always constant? I feel like planes would have different size tanks

How do plane manufacturers determine beta? I'm surprised that they have not made planes that could fly between any two points on Earth.

Beta is dependent on factors like airspeed, lift-to-drag ratio, and specific fuel consumption.

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{\text{tank}}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{\text{tank}} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{\text{plane}}}{s_{\text{bird}}} \sim \frac{\beta_{\text{plane}}}{\beta_{\text{bird}}} \frac{\mathcal{E}_{\text{plane}}}{\mathcal{E}_{\text{bird}}} \left(\frac{C_{\text{plane}}}{C_{\text{bird}}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

I remember being confused in the last section about what beta was since it wasn't clearly defined. Maybe this definition should be moved up there.

I agree. This paragraph cleared up a lot of confusion with the last section.

This was a nice clarification. I wasn't sure if I had it before but now it's pretty straightforward.

Great succinct explanation.

this is a really nice description of the equations and summary of the units

Definitely the best explaining paragraph in the section thus far.

I agree. In the last chapter I was very confused about beta and epsilon. This paragraph explains them much better. I also like that it's concise.

Thanks for explaining this so quickly!

I think ,C, is more appropriate here.

what does the drag coefficient depend on again?

It depends on the both the geometry and the material.

http://en.wikipedia.org/wiki/Drag_coefficient

can you define this above?

wow this is pretty interesting...so theoretically if a bird had a lot of energy it could fly the same distance as a plane

isn't that true of anything in the air then? this seems like a silly statement.

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{tank}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{tank} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{plane}}{s_{bird}} \sim \frac{\beta_{plane}}{\beta_{bird}} \frac{\mathcal{E}_{plane}}{\mathcal{E}_{bird}} \left(\frac{C_{plane}}{C_{bird}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

Well surely a Cessna can't fly the same distance as a Boeing 747... I feel like the assumptions made in this problem were too big, and thus you end up with a result like all planes fly the same distance. Until I see a Cessna and a 747 flying the same distance myself, I won't believe it.

This is from the wikipedia page on drag coefficient:

0.021 F-4 Phantom II (subsonic) 0.022 Learjet 24 0.024 Boeing 787 [16] 0.027 Cessna 172/182 0.027 Cessna 310 0.031 Boeing 747 0.044 F-4 Phantom II (supersonic) 0.048 F-104 Starfighter 0.095 X-15 (Not confirmed)

so you're right - the assumption that the drag coefficient is the same is not 100% correct, but they don't actually differ that much, so all planes can roughly fly the same distance.

I really want to see an answer to this. I know a single-engine two seater is not going to be making it across the ocean anytime soon, so how can they possibly go the same distance?

Different fuel tank sizes?

The point about Cessna's is a good one and shows me that my statement was too hasty. Rather, 747's and other long-distance planes, which fill a similar "ecological" niche, have about the same shape and drag coefficient and range.

But there are several other niches: (1) medium-distance commuter flights (e.g. Boston to Washington, DC) or (2) short-distance commuter (e.g. Boston to Bar Harbor, Maine or Nantucket). For those kinds of journeys, a 747 is totally unsuitable. The leading planes in each of those niches has a very different design than a 747 (and a very different range).

I agree...I really want to know what order of magnitude this is working on? just about anything can appear to be equal if you're looking at it with a large enough scale.

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{\text{tank}}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{\text{tank}} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{\text{plane}}}{s_{\text{bird}}} \sim \frac{\beta_{\text{plane}}}{\beta_{\text{bird}}} \frac{\mathcal{E}_{\text{plane}}}{\mathcal{E}_{\text{bird}}} \left(\frac{C_{\text{plane}}}{C_{\text{bird}}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

Wow I didn't know you could just cancel things out like that.

isn't everything proportional to 1? I don't get the point of this line.

It's just a way of saying that s is constant. x alpha y means that x = (some constant) * y, so x alpha 1 means that x = constant.

thanks! the explanation helped.

I don't agree with the fact that the fuel fraction is equivalent. Bigger planes carry more fuel so will be able to fly further. They may weigh more but I'd have to see data to believe it.

I'm assuming this is for planes of similar engine types? i.e. prop engines vs. turbines?

I believe that's a correct assumption as we're only looking to compare planes based on size, not the efficiency of their engines. By assuming they are geometrically similar, we also assume they are similar in engine.

Only all planes that are the exact same size and mass...which is pretty obvious already. what does this show??

....if they have the samedrag coefficient, fuel capacity and fuel efficiency

....if they have the samedrag coefficient, fuel capacity and fuel efficiency

....if they have proportional fuel capacities (B was a fuel fraction).

Upon initial glance, it looked surprising, but reading all the assumptions made, the result is pretty obvious. Keeping beta and C constant, means that even though there planes of different masses, heavier ones will naturally carry more fuel.

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{\text{tank}}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{\text{tank}} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{\text{plane}}}{s_{\text{bird}}} \sim \frac{\beta_{\text{plane}}}{\beta_{\text{bird}}} \frac{\mathcal{E}_{\text{plane}}}{\mathcal{E}_{\text{bird}}} \left(\frac{C_{\text{plane}}}{C_{\text{bird}}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

this is cool - but does engine efficiency play no part?

You can account for the efficiency in the energy density of the fuel?

I was really surprised that this is how this worked out!

But I feel like this doesn't take into account that some planes hold a LOT more people...

Yeah, I don't know... My intuition is fighting against the math here.

But I feel like the weight would still be proportional to the plane's volume. A bigger plane would hold more people than a smaller plane, but all big planes probably have similar mass.

Indeed, this does seem quite counter-intuitive. But after reviewing the Breguet range equation, this conclusion seems sound.

My intuition against this is really different than the other comments in this thread. I have a problem with all planes traveling the same distance because of the implications that would have on the airline industry. If it really takes the same amount of energy to move any sized plane, then flights would not be scheduled as they are.

I think that each airport will focus its flight schedule to maximize profit. For example you won't find any airports in Nebraska that will fly you all over the USA.

I followed, but am still confused. This, by experience, just doesn't seem very logical.

I agree, why do they use small shuttle planes to take people from new york to boston when they could just as easily use a big one?

it may be that smaller commuter flights run at more times carrying fewer people making them more efficient for the airline to run small planes rather than large empty ones. Also the larger planes may have a larger fuel fraction.

How true is that in reality?

I don't know if this is a curiosity equation or not, but I don't believe a true answer is the purpose of the statement.

$$\text{interesting quantity} \equiv \frac{\langle x^2 \rangle}{t}.$$

This quantity is so important that it is given a name – the diffusion constant – and is tabulated in handbooks of material properties.

4.1 Flight range versus size

How does the range depend on the size of the plane? Assume that all planes are geometrically similar (have the same shape) and therefore differ only in size.

Since the energy required to fly a distance s is $E \sim C^{1/2}Mgs$, a tank of fuel gives a range of

$$s \sim \frac{E_{tank}}{C^{1/2}Mg}.$$

Let β be the fuel fraction: the fraction of the plane's mass taken up by fuel. Then $M\beta$ is the fuel mass, and $M\beta\mathcal{E}$ is the energy contained in the fuel, where $\mathcal{E}$ is the energy density (energy per mass) of the fuel. With that notation, $E_{tank} \sim M\beta\mathcal{E}$ and

$$s \sim \frac{M\beta\mathcal{E}}{C^{1/2}Mg} = \frac{\beta\mathcal{E}}{C^{1/2}g}.$$

Since all planes, at least in this analysis, have the same shape, their modified drag coefficient C is also the same. And all planes face the same gravitational field strength g. So the denominator is the same for all planes. The numerator contains β and $\mathcal{E}$. Both parameters are the same for all planes. So the numerator is the same for all planes. Therefore

$$s \propto 1.$$

All planes can fly the same distance!

Even more surprising is to apply this reasoning to migrating birds. Here is the ratio of ranges:

$$\frac{s_{plane}}{s_{bird}} \sim \frac{\beta_{plane}}{\beta_{bird}} \frac{\mathcal{E}_{plane}}{\mathcal{E}_{bird}} \left(\frac{C_{plane}}{C_{bird}}\right)^{-1/2}.$$

Take the factors in turn. First, the fuel fraction β_{plane} is perhaps 0.3 or 0.4. The fuel fraction β_{bird} is probably similar: A well-fed bird having

Following your notes, this makes sense, but I guess I still have trouble how you can just simply everything like this.

This is a very easily grasped concept for symmetry.

Fuel really makes up 40% of a plane's weight? That seems high when you factor in engines, people, and luggage....

I also agree that the 30-40% percent of airplane mass is due to fuel seems rather high. I would have guessed 25% myself. But then again I would have guessed 20% for the fat percentage on a bird, so I guess it works out. However, if I had guessed 40% for plane but 20% for bird, then my calculation would approximate that a plane has the same maximum travel range as a bird. How do we avoid these errors, or are they inevitable?

25% and 30% are well within the error of our estimation.

Also, fuel is really heavy. Think about lugging one of those little tanks of gasoline. Now think about how long it takes to fuel a plane and they're probably using a pump that can provide a lot more fuel per second than a standard gas pump. if you really think about it, 35% is totall reasonable.

no it's actually totally reasonable. the fuel is a very large portion, which is in part why it doesn't really matter how many people are on board and why they didn't charge for luggage before.

I agree. Further detail on how you arrived at this estimate would be useful.

If you calculate it out: A 747 has a max fuel capacity of about 45000 gallons and @ 6lbs/gal for Jet A the fuel weighs 270,000lbs. A fully loaded aircraft is about 950,000 so the fuel is about 30% of the aircraft

I don't see how fuel fraction of the bird and the plane would be similar??? What does a fuel fraction mean precisely?

I'm a little confused about this too. Is this the overall efficiency of converting the stored energy from the fuel to kinetic energy?

Confused again on how these ratios are known. Should we just look these up or is this common knowledge?

Well, one could certainly look them up, but I suspect that divide-and-conquer would allow you to get a pretty good estimate of beta for a plane.

fed all summer is perhaps 30 or 40% fat. So $\beta_{plane}/\beta_{bird} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{plane}/\mathcal{E}_{bird} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{bird}/C_{plane} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

so the thing is, you can't use up all of a bird's fat in one go like you can with fuel without bird health serious issues. i'm guessing this is where we just blur our vision a bit but it's still a fundamental difference. if we wanted to be more precise, how would this affect our calculations?

Not the pigeons in Boston during winter...those things are OBESE.

But I doubt the pigeons in Boston are migratory birds...

how does this compare to other animals?

My question has to do with how they maintain their flight ability given they've put on weight? If people get fat, they can't run as far as they might like to.

fed all summer is perhaps 30 or 40% fat. So $\beta_{plane}/\beta_{bird} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{plane}/\mathcal{E}_{bird} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{bird}/C_{plane} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

How did you come up with those numbers? I would have guessed slightly less.

I agree I would also have guessed significantly less, but I don't have any other animals to compare these values to which would probably help.

I know for humans, normal is around 15-25%. I have no basis for applying the same to birds though.

I agree that this sounds extremely high...how would we have gotten to this otherwise?

i dont think it matters how he got there, its just trying to prove a point.

Well birds' bones are pretty small and they most use their wings for moving so I wouldn't think they have much muscle. Compared to humans 30-40% sounds about right.

So, I was kind of skeptical about this percentage (mostly just off the guess of "Well, the chicken I eat is certainly not that fatty!"), but googling around I found this book excerpt: http://tinyurl.com/ya7td8v It seems to depend on the bird, but 30 to 40% is in range.

I don't know how useful this is to you guys, but I believe it. I've had a pet bird and all they do is eat.

I also would have guessed less for both the plane and the bird, but then my ratio would still be 1 so I think it's more important to consider that they are both about the same than the exact numbers.

Well when compared to chickens, which spend a significant amount of time eating and even less time flying (I wouldn't say that they are 30% fat). Plus fat is not very dense so that means that makes it even less likely to be 30% of their weight even if there is a lot of it.

"...for humans, normal is around 15-25%"? no way, 25% seems to be a bit too much; a quarter of the body is fat? 10-15% seems more reasonable.

fed all summer is perhaps 30 or 40% fat. So $\beta_{plane}/\beta_{bird} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{plane}/\mathcal{E}_{bird} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{bird}/C_{plane} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

Also, 15-25% is totally reasonable for humans. For athletic competitive men, it's 5-10. Less than that is unhealthy, but body builders, wrestlers, rowers, etc do it for competition day. For women, it's higher due to various different bodily needs; in fact, if body fat is too low (less than 15%?), women will not menstruate.

With today's obesity epidemic, about 30% of Americans are over 35% body fat (or maybe that's 40?). Another 30% exceed 25%. Kinda eye opening, no?

That sounds surprisingly high to me also, especially since they have to be flighted, I'd think it's lower than that. Birds have SUPER low density bones and very muscular wings / diaphram so it's hard to think fat % is so high. Perhaps our force-fed chickens would be 40%.... (sorry, just watched Food Inc.)

I like the thing you said in class that the density of gasoline is the same of fat/calories. Very useful fact

This is always very interesting to me. I guess this is part of the reason fat is so hard to get rid of, it's got as much energy as jet fuel lol

haha wow, did not know this

How do we know this?

I think we can reason it because fat and jet fuel are both consumed by oxidizing long hydrocarbon chains, so it makes sense that they would release similar amounts of energy.

I find it difficult to assume that animal metabolism is about as efficient as a plane engine. My gut really doesn't have a good feel for this.

fed all summer is perhaps 30 or 40% fat. So $\beta_{plane}/\beta_{bird} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{plane}/\mathcal{E}_{bird} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{bird}/C_{plane} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

how would we know this, if you hadn't just told us?

Well we know that people are about 25% efficient as well that unless we are using an incredibly fine tuned machine 25% efficiency is pretty standard and a valid approximation

It is interesting analysis and reasoning, but as mentioned earlier, I would have trouble coming up with these numbers and approximations on my own without use of an outside resource, which would have made my resulting conclusion different.

I don't really see how to factor this 25% in... the eplison is a measure of "energy density" —> does energy density take into account engine efficiency, i just assumed it was the amount of energy in gasoline...or fat?

I think the 25% efficiency is just another check to confirm that comparing birds to planes is sane. Also, it could be factored into the energy density: if the bird had twice the efficiency, but half the fat percentage, then the effective epsilon_plane over epsilon_bird would still be one.

Also, all of these ratios seem valid to within a factor of two, so it seems pretty reasonable to hope that the extra little factors will roughly cancel (as they did for the engine efficiency and drag coefficient terms in lecture on Friday).

True, but still I suppose it works for this

if you're admitting that it's not a great approximation, why are we allowed to make it?

because this class is about approximation–since we're really stuck here, we have to make some sort of assumtion to move on in our estimation

I agree. I believe it was made simply to allow us to move forward.

Completing the calculation gives us the advantage in that we then have an answer to work with (or refine, if necessary).

Like the comparison in calculations between birds and planes

I think this part really amuses me! very interesting!!!

I agree! The explanation as well as the end result were really interesting!!!

fed all summer is perhaps 30 or 40% fat. So $\beta_{plane}/\beta_{bird} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{plane}/\mathcal{E}_{bird} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{bird}/C_{plane} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

what makes non migrating birds different?

If they do not migrate, then they most likely wouldn't have as long of a maximum range since they are not accustomed to it. Also, it would be hard to try and predict for the number since the birds never test their maximum range.

this seems like a stretch to me

What about other planes and other birds? Does it really scale across models/species?

this is a very surprising, yet cool, result.

I really want to believe all of this should always add up, but I can't help but think part of it is coincidence.

maybe other examples can be shown to show more proportionality reasoning

...Well in reality, it doesn't add up. Lots of assumptions were made that were very liberal. Most birds probably can't fly the same range as a Godwit. Most jets probably don't have the same range as a 747. That being said, I think this comparison is a really neat trick.

I agree that it's surprising. Aren't planes modeled after birds anyway?

I agree. This is awesome.

This result is quite surprising to me. I can see where it is coming from, but still I am quite surprised.

fed all summer is perhaps 30 or 40% fat. So $\beta_{plane}/\beta_{bird} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{plane}/\mathcal{E}_{bird} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{bird}/C_{plane} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

Does this mean that the godwit was 30 to 40% lighter when it arrived?? It seems implausible that an animal would lose a third of its body weight in one trip... but maybe..

That actually would not surprise me as much. Animals implicitly do things that have been done for years and years, and losing a lot of body weight is one of them. Think of penguins raising their young, or polar bears.

It also took the bird way longer than the plane. He didn't need to drink any water that whole time?

this is a good point, how do we account for the difference in velocities? also, it doesn't seem like all birds could make this long of a trip...

The difference in velocities comes from a difference in size. A bird is most energy-efficient at a lower speed, but the net distance ends up being the same (because we assumed beta, C, and epsilon were the same for each object).

I believe birds do generally stop to eat, drink, and sleep during migration. So the bird has probably not lost all of the weight it invested into the flight, since it would have gained some back during the trip.

Wouldn't this not be a nonstop flight then?

does anyone know the time that the godwit took to make this flight?

For those intrigued: http://news.nationalgeographic.com/news/2007/09/070913-longest-flight.html

The maximum range of a Cessna is something like 1400km. That's pretty far off...

That may be because a 747 has jet engines whereas a cessna is just a prop plane.

fed all summer is perhaps 30 or 40% fat. So $\beta_{plane}/\beta_{bird} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{plane}/\mathcal{E}_{bird} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{bird}/C_{plane} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

A well fed airplane:

http://www.greenpacks.org/wp-content/uploads/2008/10/record-breaker-a-bar-tailed-godwit.jpg

that is definitely not what I expected for a bird that flies so far. I was expecting something much more lanky. The godwit must make a lot of stops along the way. This is probably still a good comparison because the godwit probably doesn't eat anything while it's floating around the open ocean.

He specifically says it's the longest nonstop flight by a bird.

True. It is still a hard fact to swallow but is pretty good at supporting the approximations and calculations done so far in the reading.

It sleeps while flying!

It might be useful to include a brief tie back to how this interesting result used symmetry...i.e what irrelevant complexities were thrown out

I was wondering about that also. Was symmetry used when you took the ratio of different measuremts?

Is this relevant? Shouldn't it be just as likely that this wasn't true, especially since the amount of weight of fuel by each could be different?

How do these birds eat? Do they make any stops on the way?

Well, I guess they burn off fat but I guess the better questions is whether they need to drink water.

but is this only true for the extreme cases? are their are birds that can fly a comparable distance?

Great fun fact!!!

Now this point is a strong one. It really strikes home the point of proportionality. The plane example was a bit reliant on assumptions, but relating it to birds and using fuel proporationality method, made the point very clear.

Never would have guessed this.

fed all summer is perhaps 30 or 40% fat. So $\beta_{\text{plane}}/\beta_{\text{bird}} \sim 1$. Second, jet fuel energy density is similar to fat's energy density, and plane engines and animal metabolism are comparably efficient (about 25%). So $\mathcal{E}_{\text{plane}}/\mathcal{E}_{\text{bird}} \sim 1$. Finally, a bird has a similar shape to a plane – it is not a great approximation, but it has the virtue of simplicity. So $C_{\text{bird}}/C_{\text{plane}} \sim 1$.

Therefore, planes and well-fed, migrating birds should have the same maximum range! Let's check. The longest known nonstop flight by an animal is 11,570 km, made by a bar-tailed godwit from Alaska to New Zealand (tracked by satellite). The maximum range for a 747-400 is 13,450 km, only slightly longer than the godwit's range.

This is really cool!

That is cool.... I think this is a cool intro to proportionality and a nice example to follow it.

stupid boxes. stop popping up all the time!

I understand that it is frustrating. Please describe what actions precede the boxes popping up, so that the user-interface bug(s) can be found and fixed.

Hi, I would love to be able to help with this box problem, but in order for me to fix the bug, I need to understand where it comes from. If it happens next time, could you please let me what you were typing or clicking when the boxes started to pop up. Also, a screenshot would be great, ideally. Best, Sacha <sacha@mit.edu>

Is there anyway to access boxes that have been overlaid by multiple boxes?

yes, if you clicking multiple times on the box you're trying to select, it will eventually come to the foreground and become selected.

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

wouldn't the easiest method here be to use a spring model...? that seems to be what the image is implying.

I found printing the pdf and reading it before reading the version with the comment helped me get through the reading better because it gave me a clean look at the section. I think recommending that people make sure to read it without reading the comments while going along will help people figure out the readings better (and helped me note specific things that I thought were questionable).

I feel like these are great examples. Can we do something that is more course 6 related just to get a different frame of reference?

I really enjoyed this section. The examples were clear and to the point. Proportionality is making more sense to me than the previous topics.

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

Read Section 4.2 and Section 4.3.1 (the first part of the section on animals jumping) for the Tuesday memo.

The page breaks are not great, but I don't worry too much about that because they'll all change anyway after the revision.

go away you stupid box.

It feels like considering all mountains to be cubical is such a big approximation... Mountains are all sloped...

This sentence is intriguing! I really like this as the beginning of a new section- it makes me eager to read more!!

how high is too high? haha

Agreed, great sentence...cut and dry and right down to the point, letting the reader know exactly what will be covered.

i'm just curious whether anyone has anyone ever wondered about this?

I have not, and not only do I find this first sentence intriguing, but I have no idea where we'll begin with this calculation. Usually with this class I could at least guess an approach we could take even if I didn't know how to do it.

I guess this will be useful for our homework problem

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

I would have guessed they're more triangular.

me too, but if you assume them all to be that way it probably doesn't matter.

Wait I'm so confused - I've never seen a mountain that even vaguely resembles a cube shape.

While I've only read the next few lines, I don't think the shape of the mountain really matters. It is just an easy approximation to prove the point that mountains can't exceed a certain height.

The way that I envisioned this was a series of smaller & smaller cubes...however, rereading this made me realize that it was just my brain trying to make it make since.

My guess is any solid 3D structure would work, a cube is probably to simplify math

I think for this estimation, we're just looking at mountains as a blob of mass, to try to prove that mountains cant exceed a certain height

Well, pyramidal / tetrahedral would be the best guess, but really, a pyramid's volume is 1/3 of that of the cube (and since pyramids are an underestimate - mountains are more parabolic); the error factor is within 10ˆ0.5.

i think that's why this says "assume that all mountains are cubical," not "model all mountains as cubical.

This is definitely an interesting assumption.. Can we have some reasoning?

Although it may be misleading, it will give us a relationship. Which is always nice to have before trying to model something to actual dimensions.

The important thing to take away is that mass heightˆ3, or any linear dimensionˆ3. It assumes that all mountains have the same shape (and density), regardless of what that shape is.

This reminds me of the "assume the cow is a sphere" joke

Wouldn't a cone make more sense? or does this complicate too much?

Couldn't it be equally as easy to estimate a mountain as a right triangle?

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

In addition to solving problems like these while making assumptions like all planes are about the same and all mountains are cubical and same material, etc. it would be nice if you discussed briefly the effects of differences between the planes or mountains.

this whole paragraph is awkward....and redundant with the third paragraph

Uh, not it's not (to either of your points).

?

probably means to insert a section number here that he does not yet know.

Yeah I'm pretty sure that's what's happened.

Could be an error with some formatting or an image / equation that got lost in file format?

something to do with not oversimplifying?

is this a typo or a question posed to the read?

How can we assume that the reason that mountains aren't taller is the fact that there is too much pressure at the bottom. Could there not be some other limiting factor that we are not addressing before making this assumption?

I agree, even if I can't think of what another limiting factor could be.

Redundant, but kind of nice

Maybe rephrase the end of the sentence as "so we need to estimate the force and the area"

Probably should have read this before I tackled the pset.

Definitely. I was struggling until I gave up to do the reading.

i don't think f = mg is a hard concept to understand here.

True, but skipping this step of explanation can make the next bit more convoluted, especially for people who do not see equations on a day to day basis.

Yeah – even if it's a simple equation, I don't think having it there detracts from anything.

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

I don't really understand how F being proportional to mass makes F/m independent of the mass. Isn't mass a part of F/m?

So when we say that some ratio is "independent" of a quantity, are we saying that it stays constant as something changes? I don't know why I'm not understanding this intuitively.

Yes, x independent of y if changing y does not change x.

or the constant, gravity.

so the force inhibiting mountain height is independent of gravity? In this week's PSET there's a question that insinuates the height of mountains on different planets is related to gravity?

Is it possible that gravity figures into the force, such that that F/m is independent, but F is not?

yeah, when he says that F/m is independent of mass, he means it's a constant–the constant of gravity. so, F does depend on gravity, and m doesnt; look at weight: F=mg, mg/m = g. a constant. that's all he's saying here.

I know this language of "being independent" was used in the previous reading, but it's still really confusing. I feel like I shouldn't have to think this hard to figure out what the original definition of this term means.

When you say that F/m is independent of mass, it makes me think that as m changes, F/m won't change, which is clearly not true. I think another phrase might be more useful and create less confusion.

wait as m changes, F/m DOESN'T change. F=mg so as m changes F changes proportionally and thus F/m does not change. so your assertion that this statement is not true is not a proper assertion

Yeah, "When you say that F/m is independent of mass, it makes me think that as m changes, F/m won't change, which is clearly not true. I think another phrase might be more useful and create less confusion."

It is clearly true, which is why it is independent.

how is F/m independent of mass- if m is mass??

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

I dont understand how F/M is independent of mass when mass is included in it. what exactly does this mean?

It means that the ratio between force and mass does not depend on mass. For example, no matter what my mass, my force of gravity divided by my mass is always the same.

A bit of a subject jump here. It almost seemed like the l^3 relation was related to the mass relation, which made me read it twice. Looking at the proportionality below it made everything clear though.

I fell like this while part is confusingly worded & ^3. therefore F proportional to l^3

I agree with this sentiment. I read this without seeing any comments first (printed it out), and after getting through the first page had a "What?" moment. I makes sense going back to it, but just having F prop m, and m prop vol, and f prop l^3 just felt like running down an actual mountain,.

why are you using the   symbol here?

Yeah I'm confused about that as well - mass is proportional to volume, but volume is definitely equal to L^3.

maybe since we're acknowledging that we made a big assumption when we said mountains are cubical?

I had been reading   as "is approximated by" and the proportionality symbol as just that (throwing out constants, arbitrary or not).

Is this related to the proportionality with density?

I think I get why when you make the equations a fraction, they become independent, but if you could expand a little bit on the explanation, it'd be more helpfull.

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

I don't understand why this is independent of l, a little more explanation could be used.

same reason that F/m is independent of m, it takes into account m...the proportionality is unchanged

if you think about it, m=density*volume, so m/volume=density. here we're assuming that density of mountains is the same, or "constant" so the ratio of mass/volume is unchanging for different mountains.

l^3 = volume.

This was a little confusing to understand at first and I made some mistakes on the homework in this area but now that I understand it it is incredibly useful

the constant, density

This was a little confusing

I think another sentence to explain/clarify the use of constants would be helpful.

I'm still a little confused about this, you don't really explain why the equation becomes independent of l

Every time i see the word independent I think they are actually dependent on each other.... confused.

Yeah, I think I'm also confused... why is it m/l^3 is independent of l and not, say, m? Is it because m is a function of l? And that m/l is a constant...?

relating it back to density (see above comment) made it clearer for me... because i know that objects of varying sizes can have the same density

I'm confused with this wording.

Yes, we do seem to be spending a lot of words to argue a fairly straightforward point.

yeah i think the wording of this sentence is weird, and kind of hard to follow.

i see you're using parallel structure to mimic what you argued above but i don't think it's necessary again. i feel like the second time it's a "duh" point

Maybe writing out the full equation and then showing which variable s are eliminated because they are constant would make this more clear.

All he's saying is F m l^3, so F l^3.

4.2 Mountain heights

The next example of proportional reasoning explains why mountains cannot become too high. Assume that all mountains are cubical and made of the same material. Making that assumption discards actual complexity, the topic of ??. However, it is a useful approximation.

To see what happens if a mountain gets too large, estimate the pressure at the base of the mountain. Pressure is force divided by area, so estimate the force and the area.

The area is the easier estimate. With the approximation that all mountains are cubical and made of the same kind of rock, the only parameter distinguishing one mountain from another is its side length l. The area of the base is then l^2.

Next estimate the force. It is proportional to the mass:

$$F \propto m.$$

In other words, F/m is independent of mass, and that independence is why the proportionality $F \propto m$ is useful. The mass is proportional to l^3:

$$m \propto \text{volume} \sim l^3.$$

In other words, m/l^3 is independent of l; this independence is why the proportionality $m \propto l^3$ is useful. Therefore

$$F \propto l^3.$$

It would be interesting to see if any of the mountains on earth are anywhere close to this limit.

It would be interesting to see if any of the mountains on earth are anywhere close to this limit.

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

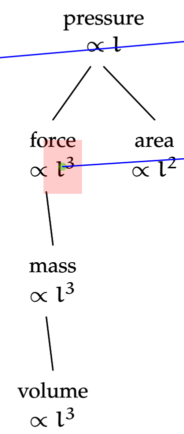

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

beautiful...I actually get it!

Yeah this was a really awesome example!

I'd change this to m

I think this works as l^3 since the pressure=force/area assertion earlier clearly combines with this to show that force/area=l^3/l^2=l

I didn't even realize he used L^3. But he didn't make an accidental mistake. He is doing the same thing he did earlier when he said mass is proportional to volume (L^3). In fact, force is as well, since the force component of an object's pressure on the ground is its mass times the little g (gravitational constant). We don't worry about "g" so the m in F=ma reduces to L^3.

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

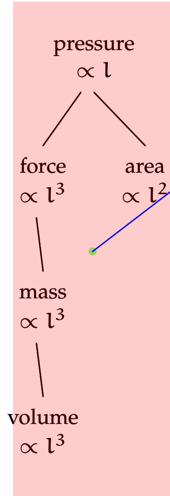

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

this doesn't make any sense to me. it's not our traditional tree...

A small introduction to this structure would be helpful.

Agreed it's new, but it's similar to what we did before. Instead of having numbers, we just keep the variables to see how things scale wrt proportion. For example, before when we were estimating mass a common tree structure was to first estimate volume and density and multiply the two. In that scenario we used numbers. In this example it's enough for us to know that they scale in proportion. In other words, mass is proportional to density and volume. (increase volume –> increase mass and increase density –> increase mass)

It's pretty similar to what we had done previously. In a tree, you stop at a branch if you know how to estimate that variable. In this case we break pressure into area and force. We can easily find area but we can't estimate the force of a mountain so we need another branch. We know that F prop to mass but we can't estimate the mass of a mountain either so we extend it to volume which is something that we know how to find

yeah, I agree–this is pretty similar to what we've seen. in the first case, there's two values we need to estimate pressure: area and force. We can estimate area right off the bat, so we use more branches for force. then we say force depends on m*g, but g is constant. so look at m. but we cant confidently guess m, so we look at m: it's density*volume. but density of all mountains is the same, so we then have volume, which using our cube approximation, we can find

I would definitely rather have the tree than not have it at all. It is always hepful to have a diagram that summarizes your steps to solving the problem. Not to mention that it is very easy to follow. (much easier than reading the entire paragraph)

i would argue for keeping the tree. though it may not be helpful for everyone, i personally (and i know that i am not alone) really like graphical approaches. those who don't like could always opt to not use it.

I think its easy enough to see how this tree is like the divide and conquer we did.

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

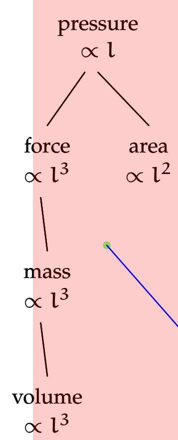

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

The reply above give a nice explanation of the tree. I only understood the first section after seeing the diagram. You should fuse the two and use the text to explain the tree instead of using them seperately.

Keep the tree. But if there was some way to put it on the previous page where you go through all the steps, that would be more helpful

I actually like this proportionality tree A LOT. Lining up the proportionalities makes this very clear.

The proportionalities are made clear but I had to read the text to understand the purpose and function of the nodes. Why not add a leaf for density and g?

i'm having a hard time understanding how to read these trees

It doesn't say anything explicitly about a maximum though, just that the pressure consistently increases with height. But perhaps the pressure can go almost infinitely high or the proportionality constant is infinitely small?

It makes sense to me that the rock would become liquid but isn't it dependent on the fact that you are assuming the mountain is a cube?

The idea here was to explain why there _is_ a maximum height, not to show what it is.

but that would be fun/interesting!

I am sure that finding it is as easy as finding the pressure (at a constant temperature) where the rock turns into liquid and finding [L] in terms of pressure to find its max height.

I don't think he's saying that the rock turns to liquid but that pressure from the height of the mountain will cause it to compress and make the base widen, like a liquid would. I could be way off here though.

I like this paragraph. It makes sense reading it.

This explanation that the rock will flow like liquid should be earlier in the section. Otherwise the reader is imagining how a mountain's height can be limited.

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

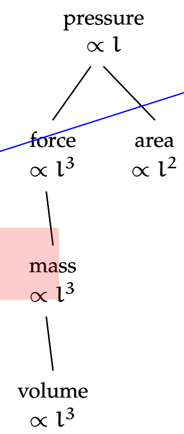

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

This is only partially true. What about erosion from rain, snow and glaciers? Mountains on mars are much larger than mountains on earth because they lack these.

Well, you can also think about it in terms of why a mountain of height very tall h couldn't even form in the first place, right...?

But to illustrate the example at hand none of this information is really relevant. It would be an entirely different problem to try and approximate how much height a mountain loses each year due to errosion.

I still think it's good to keep in mind that on earth other effects dominant the restriction on height of mountains so that you have a sanity check for your estimation. The maximum height of a mountain based on the strength of rocks should be much larger than the observable height of mountains on earth.

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

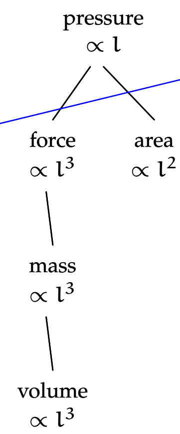

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

So the inside of the mountain turns to lava? Does the mountain shrink down into the earth at that point?

shrink? no. i think you mean sink. and sink into what? more rock. that doesnt make any sense either.

maybe the rock "overflows" the mountain and just starts rolling down, or something. perhaps like a liquid overflowing its container and spilling out the sides.

a little earth 101: reference image (http://www.nasa.gov/images/content/103949main_earth

The outer part of the earth is the "crust" ... the hard earthy parts that we know and love, yet take up less than 1% of the earth's volume.

Below the crust is the mantle ... esentilly magma, or liquid rock.

Then there's the outter core (liquid metals) & the inner core (solid metals)...both mostly iron.

...

The idea is that there is only so high that the crust can get before the mountain's weight applies enough pressure to melt the bottom of it into the mantle.

I'm sorry, I can't let this slide - the mantle is not a liquid!!!!!!! The asthenosphere (right under the crust) is hot and acts like a fluid – in fact, the mantle solid-state convects (think about it), but it is certainly not a liquid. In some place, yes, the mantle may have some melt, but not more than a few percent. And...it has been suggested that mountains (or dense bottom parts of continents) loose their roots...the dense bottoms drip down...making more melt! anyways, the point is...the mantle is not a liquid.

He said the rock flows LIKE a liquid... he didn't say it turned into liquid.

I'm guessing that's you steph

So I guess this makes sense since there's conservation of mass, because the earth could be thought of a big mountain (different shape) but with the middle turning to lava because of the pressure....

Are you implying that if the earth gets too big, there will be too much pressure and the core will melt?

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

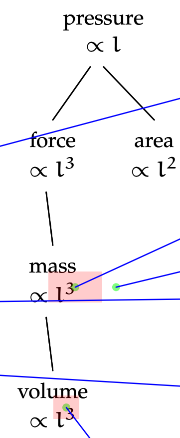

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

I'm sure this is heavily dependent on the material that the mountain is made of. Volcanic rock like granite is a lot more robust than dirt hills or mountains made of sedimentary rock

and this to volume

This actually made things alot clearer. I guess visual effects can do that to you.

this logic is a little fuzzy to me

this clears up my earlier comment about how do we know that pressure is the limiting factor. Pressure is something we know that will limit the height and lets us know that the mountain does have a max height.

Is it really necessary to even use math to show there exists a maximum height? Can't an argument just be made without all this?

Not necessarily. If we didn't know the math, we might think that the pressure on the base of the mountain was constant with increasing size, which would not predict an upper limit on height. I admit that the fact we've proven is not very difficult to grasp intuitively, but as an instructional example this is very useful, especially since we can use our intuition to confirm the result of the math.

and this to = l^3

but now modeling the mountain as a cube, not a cone or pyramid, could put us off by a factor of few right from the start, no?

Oops, next paragraph addresses this.

I like this section, and my comment is just that I feel a little jipped that I don't get the answer now! But it's understandable...

I was just about to write the same thing.

I agree as well, but it's a good thing that reading the book makes us want to read more of the book.

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

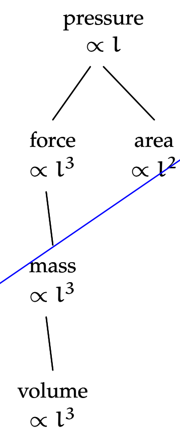

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length.* Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

this is something that meche's would just look up in a book, or run an experiment on. experiments are how the strengths are actually found.

As a course 2 student, I agree, but I think the idea is to use the knowledge you know about things like this in order to approximate and not get an exact answer.

it's one of those examples though that you never really bother to examine until someone points it out. i think it would be rather cool to derive the answer using estimation!

I think most Mech-E's probably learned this in thermo. Remember dimensional analysis?

yea i'm a mechE too but do i really EVER feel like figuring out the strength of a rock or ever expect to understand anything from the # i pick up? i'd rather approximate

Furthermore, its good to have a general intuition about these things so that you know that the numbers you look up are good or are in the correct units. Sometimes you are in the field or it may be too inconvenient or even impossible to look up particular values.

Assuming that objects take on square or cube-like shapes seems to be popular in this course. I guess it has to deal with the fact that those are relatively easy to approximate.

answer to the question someone had earlier.

It's interesting that the proportionalities make it boil down to just one factor of length. However, we've linked the base length and height. If those aren't related, you have mountains with the base not proportional to height (Dolomites, some Chinese mountain ranges), which might experience different limits.

This paragraph might more useful towards the beginning, especially considering how many people commented on the seemingly questionable estimation of a cube.

It would have been nice at the beginning to see something acknowledging that a cubical mountain seems like an odd assumption but promising to address that at the end (once we've gone through the proportionality stuff)

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

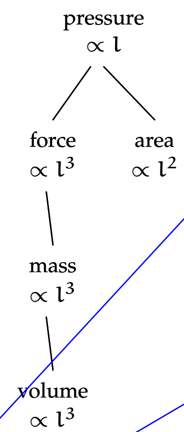

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

This isn't totally clear to me.

I think it's saying that if there is any sort of multiplier, because it is a ratio that we are looking for, the multiplier will just cancel out.

I don't think its so much a matter of the multipler cancelling out as it is how the multiplier doesn't affect the proportionality. The exact result may vary b/c of the multiplier but the general relationship is the same regardless.

ar-huh here's why, I was confused at the beginning and now i got it

So when making a proportional argument, do we always throw out constant factors? If we have a large constant, can we assume that our variables are sufficiently large to make the constant irrelevant for the sake of proportionality?

Yeah, I don't quite get this either. Why doesn't the one-third matter?

I think it has more to do with the right order of magnitude and factors like 2 or 3 don't matter too much. Of course, if you throw out too many factors eventually they add up to something more significant.

Why can't this argument be use to say there there is basically no limit to mountain height, since as the ratio of the length of the base to the height decreases the maximum height increases. So a mountain with an arbitrarily large base could be arbitrarily tall.

I really like how this paragraph questions the initial assumption of using a cube to approximate the volume of the mountain, and that it shows that even if we had used a more realistic shape, we would have reached the same conclusion.

I don't think this was too complicated to just include earlier.

But not including it demonstrated that we got the right answer without using that detail. Next time it may indeed save us a lot of effort to assume the object in question is cubical.

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

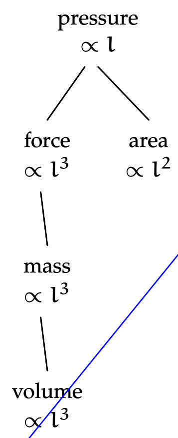

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

Does anyone actually know the answer, and are current mountains not nearly tall enough to come close to the maximum hieght?

I tried to find it on the internet or wikipedia. I'd be interested to know...

each rock type has different properties, so it's one of those things where _every_ region has a different 'maximum size.' based on the specific minerals that are in (and under) them, where in different plates that they are, and everything else. I think that every mountain is at it's maximum height.

Naw...the maximum height can only be considered when the mountain is actively being made (aka the Himalayas). Older mountains are not at their max height...erosion works way too fast! The Appalachian Mountains were once as tall as the Himalayan mountains...but are they now? Mountains only last for on the order of 10 - 100s of millions of years.

I don't know if this is true, but I once heard that the Himalayas stay at around the same height since the erosion works about as fast as the fact it's growing, which makes it sound like the maximum height is limited by erosion... is this untrue?

The Himalayas is still growing, but at a very slow rate due to erosion.

Interesting, thanks for the comments guys.

Interesting, thanks for the comments guys.

Is there a reason why you explain the cubical approximation later in the page? I guess it kind of has a nice build up effect, but people would question in the beginning why you say cubical.

So what is the maximum height of a mountain?

This make sense since scaled computations should be good comparison.

Why do you need to assume all mountains are made of the same material? Couldn't you still calculate the maximum height for every mountain based on the different rock strengths?

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

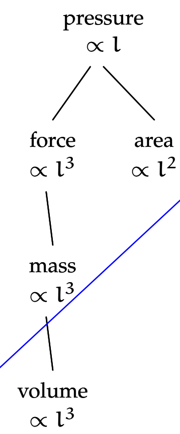

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

I don't really feel like the mountain and the jumping tie in together very well. Maybe there should be some better sort of transition here.

And the mountain doesn't transition from the flight section earlier.

But I don't think that is the point. They are just different examples showing proportional reasoning.

Sometimes transitions can be helpful though as they remind you of things you learned in the previous section that are pertinent to the next example. Sometimes we are presented with so many facts and methods in the previous section that it is a little unclear what our starting context is for the next example.

Maybe the kangaroos are planning to jump over the mountain.

proportionally?

Yeah I would like this to be clarified- I assume like the bird and 747 example it would be proportionally

I thought it was just the exact height, thought considering we are leaning proportional reasoning, it can get confusing.

I thought the 747/bird example wasn't proportional. We concluded that they could fly the same distance but the amount of fuel they need is proportional. This seems different because the kangaroo and flea aren't going to jump the same height even if they have the same muscle/mass ratio (I think.... maybe that's totally wrong).

Yeah, clarification is needed about whether or not this means proportionally...

I mean, I think this is an unnecessary comment given the examples involved. Obviously, a Meter+ size animal will jump higher than a mm size insect.

penguins can jump 6 feet.

Is this really true? I would be amazed if this were a fact considering they have tiny legs and are pretty awkward on land

The force and area results show that the pressure is proportional to l:

$$p \sim \frac{F}{A} \propto \frac{l^3}{l^2} = l.$$

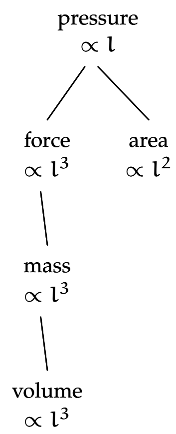

With a large-enough mountain, the pressure is larger than the maximum pressure that the rock can withstand. Then the rock flows like a liquid, and the mountain cannot grow taller.

This estimate shows only that there is a maximum height but it does not compute the maximum height. To do that next step requires estimating the strength of rock. Later in this book when we estimate the strength of materials, I revisit this example.

This estimate might look dubious also because of the assumption that mountains are cubical. Who has seen a cubical mountain? Try a reasonable alternative, that mountains are pyramidal with a square base of side l and a height l, having a 45° slope. Then the volume is $l^3/3$ instead of l^3 *but the factor of one-third does not affect the proportionality between force and length*. Because of the factor of one-third, the maximum height will be higher for a pyramidal mountain than for a cubical mountain. However, there is again a maximum size (and height) of a mountain. In general, the argument for a maximum height requires only that all mountains are similar – are scaled versions of each other – and does not depend on the shape of the mountain.

4.3 Jumping high

We next use proportional reasoning to understand how high animals jump, as a function of their size. Do kangaroos jump higher than fleas? We study a jump from standing (or from rest, for animals that do not stand); a running jump depends on different physics. This problem looks underspecified. The height depends on how much muscle an animal has, how efficient the muscles are, what the animal's shape is, and much else. The first subsection introduces a simple model of jumping, and the second refines the model to consider physical effects neglected in the crude approximations.

the way this sentence is introduced sounds way random

yeah what does this mean? it sounds like a comment we might make

a colon or semicolon instead of a period after 'underspecified' would probably help the sentence feel more connected to the next phrase, which explains what it meant.

This section is good because it prepares us well for what we'll have coming. In a lot of the past readings I've been confused about why we were making the calculations that we made.

The height of the jump (it makes the sentence clearer)

4.3.1 Simple model

We want to determine only how jump height varies with body mass. Even this problem looks difficult; the height still depends on muscle efficiency, and so on. Let's see how far we get by just plowing along, and using symbols for the unknown quantities. Maybe all the unknowns cancel.

We want an equation for the height h in the form $h \sim m^{\beta}$, where m is the animal's mass and β is the so-called scaling exponent.

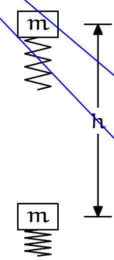

Jumping requires energy, which must be provided by muscles. This first, simplest model equates the required energy to the energy supplied by the animal's muscles.

The required energy is the easier estimation: An animal of mass m jumping to a height h requires an energy $E_{jump} \propto mh$. Because all animals feel the same gravity, this relation does not contain the gravitational acceleration g. You could include it in the equation, but it would just carry through the equations like unused baggage on a trip.

The available energy is the harder estimation. To find it, divide and conquer. It is the product of the muscle mass and of the energy per mass (the energy density) stored in muscle.

To approximate the muscle mass, assume that a fixed fraction of an animals mass is muscle, i.e. that this fraction is the same for all animals. If α is the fraction, then

$$m_{muscle} \sim \alpha m$$

or, as a proportionality,

$$m_{muscle} \propto m,$$

where the last step uses the assumption that all animals have the same α.

For the energy per mass, assume again that all muscle tissues are the same: that they store the same energy per mass. If this energy per mass is $\mathcal{E}$, then the available energy is

$$E_{avail} \sim \mathcal{E} m_{muscle}$$

or, as a proportionality,

Actually quite clever, since this way, we can disregard actual systems of mechanical movement between species.

I am not convinced that we can just use body mass. There was an article in the BBC about an ant that could lift 100 times its mass.

http://news.bbc.co.uk/2/hi/uk_news/8526086.stm

What is this "so on" referring to?

the factors mentioned earlier, such as: how much muscle the animal has, what the animal's shape is, etc.

These are factors that were mentioned before. Other factors I can think of include how athletic a person is, how much a person has practiced, etc.

I don't think I'm reassured by the same blind hope that you are. Then again, perhaps it's not so blind for you.

this seems like a very tricky statement. unknowns, by definition, can't cancel. They'll just change the values by unknown scalars.

The point is that they are unknown quantities but if we can approximate the level of them we might be able to have them equal one another.

I see what you're saying... However, i think he is referring to the possibility of having the same variable somewhere in the numerator and also somewhere in the denominator. For example, you could have height in the numerator and volume in the denominator. But volume uses height, so the height cancels.

This statement just seems out of the place the way that it's worded.

It kind of makes the paragraph sound like you have the problem all worked out - know the variables are going to cancel - but are just playing along for the reader. It sounds awkward.

I read this statement to mean that the factors contributing to the max height would cancel..for ex., higher body mass would offset more muscle. Not necessarily would the units cancel, just their relation to whether it enabled the animal to jump higher or not.

I sometimes get frustrated when things just cancel. But it is interesting to just try without much information than give up, because I think ordinarily I would just give up.

4.3.1 Simple model

We want to determine only how jump height varies with body mass. Even this problem looks difficult; the height still depends on muscle efficiency, and so on. Let's see how far we get by just plowing along, and using symbols for the unknown quantities. Maybe all the unknowns cancel.

We want an equation for the height h in the form $h \sim m^{\beta}$, where m is the animal's mass and β is the so-called scaling exponent.

Jumping requires energy, which must be provided by muscles. This first, simplest model equates the required energy to the energy supplied by the animal's muscles.

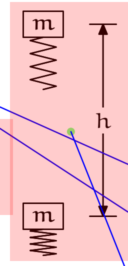

The required energy is the easier estimation: An animal of mass m jumping to a height h requires an energy $E_{\text{jump}} \propto mh$. Because all animals feel the same gravity, this relation does not contain the gravitational acceleration g. You could include it in the equation, but it would just carry through the equations like unused baggage on a trip.

The available energy is the harder estimation. To find it, divide and conquer. It is the product of the muscle mass and of the energy per mass (the energy density) stored in muscle.

To approximate the muscle mass, assume that a fixed fraction of an animals mass is muscle, i.e. that this fraction is the same for all animals. If α is the fraction, then

$$m_{\text{muscle}} \sim \alpha m$$

or, as a proportionality,

$$m_{\text{muscle}} \propto m,$$

where the last step uses the assumption that all animals have the same α.

For the energy per mass, assume again that all muscle tissues are the same: that they store the same energy per mass. If this energy per mass is $\mathcal{E}$, then the available energy is

$$E_{\text{avail}} \sim \mathcal{E} m_{\text{muscle}}$$

or, as a proportionality,

If I was just doing a problem I haven't seen before, this method makes me feel very uneasy. But maybe by being very consistent with what does unknowns are, it could work out. I feel like a lot of problems would arise by not being consistent with when the unknown variables should appear and when they should not.

How do you know you want the exponential relationship. I think there was a "jump" in logic.

Can you motivate the need for /beta now?

Well we don't know if h will be proportional to m, m^2, m^3.. etc

Yep, beta is an exponent, not a constant in front of m.

This is just what we did in class today with n^k for the schooling method.

You should specify whether this is the animals height, or the height of the jump. could be confusing

i'm confused.

I am too but maybe I can be a little more specific. How do you know that you are looking for a result in this form? There may be a simple step I am missing.

I just read the comment above and am no longer confused this makes perfect sense

agreed, but i still have no idea how to proceed from here. i realize that below will explain, but i would not have been able to do this alone

I don't know that the point is to be able to do this alone on your first try seeing it but to show you an example so that next time you can do it alone

Could we model these as springs? It wasn't mentioned, but this drawing gives me the impression of using 1/2kx^2

Doesn't this assume all animals have the same center of mass? An animal with little mass in their legs could make a small jump and pull up their legs and clear a height larger than the change in the position of their center of mass.

I really like the order that we approach things in this section. It's very logical and intuitive, and even reproducible.

4.3.1 Simple model

We want to determine only how jump height varies with body mass. Even this problem looks difficult; the height still depends on muscle efficiency, and so on. Let's see how far we get by just plowing along, and using symbols for the unknown quantities. Maybe all the unknowns cancel.

We want an equation for the height h in the form $h \sim m^{\beta}$, where m is the animal's mass and β is the so-called scaling exponent.

Jumping requires energy, which must be provided by muscles. This first, simplest model equates the required energy to the energy supplied by the animal's muscles.

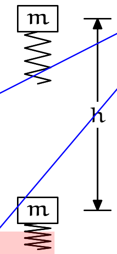

The required energy is the easier estimation: An animal of mass m jumping to a height h requires an energy $E_{jump} \propto mh$. Because all animals feel the same gravity, this relation does not contain the gravitational acceleration g. You could include it in the equation, but it would just carry through the equations like unused baggage on a trip.

The available energy is the harder estimation. To find it, divide and conquer. It is the product of the muscle mass and of the energy per mass (the energy density) stored in muscle.

To approximate the muscle mass, assume that a fixed fraction of an animals mass is muscle, i.e. that this fraction is the same for all animals. If α is the fraction, then

$$m_{muscle} \sim \alpha m$$

or, as a proportionality,

$$m_{muscle} \propto m,$$

where the last step uses the assumption that all animals have the same α.

For the energy per mass, assume again that all muscle tissues are the same: that they store the same energy per mass. If this energy per mass is $\mathcal{E}$, then the available energy is

$$E_{avail} \sim \mathcal{E} m_{muscle}$$

or, as a proportionality,

I think this explanation should have come much early in the proportionality concept.

Good analogy that brings home a very good idea when doing proportionality problems: since everything is relative to each other, anything in common can be replaced with a 1.

but wont you need it when it comes to actual calculations?

This was hilarious!

I also liked the comparison.

The humor lightens up the material and makes it more fun to read :) good job!

Agreed. I enjoyed this as well.

The humor is good.

Definitely a great analogy. I find my self doing both, carrying unnecessary baggage on trips, and unnecessary quantities in equations all the time, never thought of them as the same mistake.

is there a reason why we're going through it here?

I agree, it seems like a strange time to bring up how much energy is necessary.

Because it was mentioned as proportional to mh, thus if we can solve for m (we can assume the mass of the animal) and we can solve for E, we can get h

I am confused about what available energy is. If we already know how to calculate the amount of energy it takes to jump, why do we need available energy?

I liked this. I should really learn to think in terms of energy densities more when making approximations.

so out of curiosity, what is the power density in human muscle compared to say gas, or lithium-ion batteries?

4.3.1 Simple model

We want to determine only how jump height varies with body mass. Even this problem looks difficult; the height still depends on muscle efficiency, and so on. Let's see how far we get by just plowing along, and using symbols for the unknown quantities. Maybe all the unknowns cancel.

We want an equation for the height h in the form $h \sim m^{\beta}$, where m is the animal's mass and β is the so-called scaling exponent.

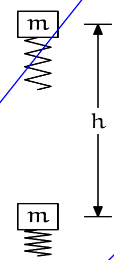

Jumping requires energy, which must be provided by muscles. This first, simplest model equates the required energy to the energy supplied by the animal's muscles.

The required energy is the easier estimation: An animal of mass m jumping to a height h requires an energy $E_{jump} \propto mh$. Because all animals feel the same gravity, this relation does not contain the gravitational acceleration g. You could include it in the equation, but it would just carry through the equations like unused baggage on a trip.

The available energy is the harder estimation. To find it, divide and conquer. It is the product of the muscle mass and of the energy per mass (the energy density) stored in muscle.

To approximate the muscle mass, assume that a fixed fraction of an animals mass is muscle, i.e. that this fraction is the same for all animals. If α is the fraction, then

$$m_{muscle} \sim \alpha m$$

or, as a proportionality,

$$m_{muscle} \propto m,$$

where the last step uses the assumption that all animals have the same α.

For the energy per mass, assume again that all muscle tissues are the same: that they store the same energy per mass. If this energy per mass is $\mathcal{E}$, then the available energy is

$$E_{avail} \sim \mathcal{E} m_{muscle}$$

or, as a proportionality,

is this a valid assumption here? is it a necessary one?

I think it's pretty legit, and it's definitely needed, since right above, we estimated energy per mass in muscle, so to estimate energy we need to estimate the muscle mass in an animal, and we're assuming that for most animals, this fraction is probably pretty similar

I wouldn't have made this assumption in my first pass through. Its tough to accept, knowing that it is not true, but only roughly true. Knowing that, aren't we limited to making the assumption that all animals jump "roughly" the same height?

isnt this already detracting from the original question? of which animals jump hihger? if they all have the same muscle fraction, and then i assume we'll say they all have the same muscle efficiency, we'll just get that they all jump the same

No, even if they have the same muscle fraction and efficiency, their masses are different.

But we're trying to differentiate between animals. This groups them all together. How can we possibly tell which animals jump higher, if they're all the same?

We're saying that their muscle fractions may be the same, but their masses are still different. If our final answer depends on mass, then a flea and a kangaroo will not have the same jump height.

i don't like the use of alpha as the fraction because this is used for proportionality

Yeah it is a little confusing at first...

You might want to use a greek letter other than alpha here. Alpha looks a lot like the "proportional to" symbol and might confuse people.

that is exactly what i was thinking

definitely. there's so many greek letters available to be used too...

4.3.1 Simple model

We want to determine only how jump height varies with body mass. Even this problem looks difficult; the height still depends on muscle efficiency, and so on. Let's see how far we get by just plowing along, and using symbols for the unknown quantities. Maybe all the unknowns cancel.

We want an equation for the height h in the form $h \sim m^{\beta}$, where m is the animal's mass and β is the so-called scaling exponent.

Jumping requires energy, which must be provided by muscles. This first, simplest model equates the required energy to the energy supplied by the animal's muscles.

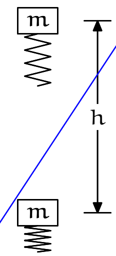

The required energy is the easier estimation: An animal of mass m jumping to a height h requires an energy $E_{jump} \propto mh$. Because all animals feel the same gravity, this relation does not contain the gravitational acceleration g. You could include it in the equation, but it would just carry through the equations like unused baggage on a trip.

The available energy is the harder estimation. To find it, divide and conquer. It is the product of the muscle mass and of the energy per mass (the energy density) stored in muscle.

To approximate the muscle mass, assume that a fixed fraction of an animals mass is muscle, i.e. that this fraction is the same for all animals. If α is the fraction, then

$$m_{muscle} \sim \alpha m$$

or, as a proportionality,

$$m_{muscle} \propto m,$$

where the last step uses the assumption that all animals have the same α.

For the energy per mass, assume again that all muscle tissues are the same: that they store the same energy per mass. If this energy per mass is $\mathcal{E}$, then the available energy is

$$E_{avail} \sim \mathcal{E} m_{muscle}$$

or, as a proportionality,

Does it matter that all animals have the same alpha for this to work? I thought the proportionality only meant that they all varied linearly (similarly) with overall mass, but not necessarily with a particular constant associated with all of them (because then you could just state that m_muscle = alpha*muscle, right?)

He does state earlier that m_muscle= alpha * m but I think for approximation's sake it makes sense to assume that alpha is the same for all animals because then it will just cancel out in our ratios.

I fell like you have to do this but it is interesting to think about whether this is true when comparing fleas and kangaroos since they have such different skeletal structure.

I'm not sure that makes sense, some animals work a lot harder than others, or have more effecient systems! oh well

i agree, but how else are we going to move on...

Exactly, it's better to go forward incorrectly than stay stuck correctly.

This is quite an assumption, there are many other things that come together to make up the mass of an animal (i.e. bones, organs).

right here he's only talking about muscle–assume all muscle tissue is the same. we know it's not, that's why it's an assumption

I guess with both this point and the point about alpha above I would let this slide if I wasn't thinking of, say, a flea versus a kangaroo because it's difficult for me to believe that their body functions similarly at all in terms of muscle, muscle mass, etc.

$$E_{avail} \propto m_{muscle},$$

where this last step uses the assumption that all muscle has the same energy density $\mathcal{E}$.

Here is a tree that summarizes this model:

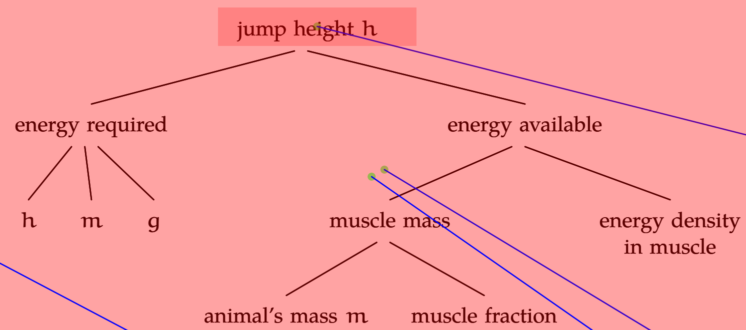

Now finish propagating toward the root. The available energy is

$$E_{avail} \propto m.$$

So an animal with three times the mass of another animal can store roughly three times the energy in its muscles, according to this simple model.

Now compare the available and required energies to find how the jump height as a function of mass. The available energy is

$$E_{avail} \propto m$$

and the required energy is

$$E_{required} \propto mh.$$

Equate these energies, which is an application of conservation of energy. Then $mh \propto m$ or

$$h \propto m^0.$$

In other words, all animals jump to the same height.

Comments on page 4

I feel like we've been making so many assumptions, such as the fraction of muscle is fixed, the efficiency is the same across all animals, but is there a way to account for these factors at the end? or they are not significant in our calculations. I think these factors can contribute to a lot of variations, unless they all got canceled out at the end

I think it's valid since we're only looking for a ballpark number in the end. It's highly unlikely that one animal will be an order of magnitude more efficient than another animal. The muscle efficiency assumption is also reasonable, as muscle is all essentially the same tissue type, and it functions as a spring.

I feel like the surface they're jumping from is important too for energy loss reasons. I guess you can wrap it all up in the energy argument, but it might be something to mention, if merely as an aside.

Hmm, interesting point. Surely assuming it to be an elastic reaction isn't too far of a stretch?

wow, yeah that would probably make a difference...didn't think of that

Since we're doing a comparison, we can assume they are jumping on the same surface. So, really, it doesn't matter. It's not like one is jumping off concrete and another rubber. It's either both concrete or both rubber.

hmmm any body else feel like they're going through case in point lol

I like how this tree combines proportions with divide and conquer. It makes so much sense!

I really like how the trees have carried throughout the notes to help tie everything together. I'm a visual learner to some degree, and they greatly help with clarification.

This diagram really cleared things up for me

$$E_{\text{avail}} \propto m_{\text{muscle}},$$

where this last step uses the assumption that all muscle has the same energy density $\mathcal{E}$.

Here is a tree that summarizes this model:

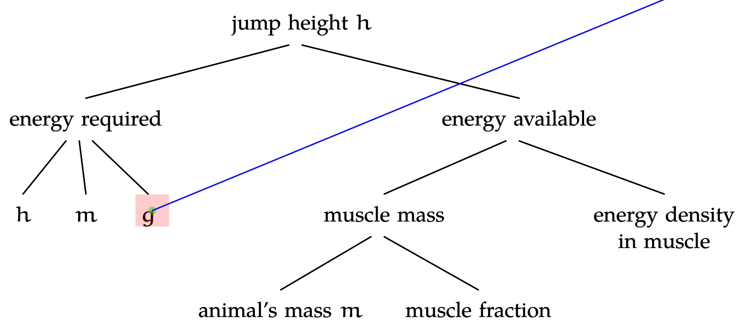

Now finish propagating toward the root. The available energy is

$$E_{\text{avail}} \propto m.$$

So an animal with three times the mass of another animal can store roughly three times the energy in its muscles, according to this simple model.

Now compare the available and required energies to find how the jump height as a function of mass. The available energy is

$$E_{\text{avail}} \propto m$$

and the required energy is

$$E_{\text{required}} \propto mh.$$

Equate these energies, which is an application of conservation of energy. Then $mh \propto m$ or

$$h \propto m^0.$$

In other words, all animals jump to the same height.

why are we including this here?

I thought 'g' was supposed to be ignored...

I think he's putting 'g' there in order to show that he is not forgetting that it exists in this problem. Even though it was mentioned before that all animals experience the same g so we can ignore it, it's still good practice to have it here in the tree diagram in order to make it cognizant to readers that it the energy required for a jump depends on g.

yeah, but why didn't he include it in the earlier problem when we said that F is proportional to m–here E is proporitonal to h*m...same kind of thing right?

i think it's because the equation for E is mgh, and all three variables feed into the equation for E. even if they cancel later, it still has to feed in now.

there should just be something in parenthesis or some other way of indicating that we won't be using it to avoid confusion and stay consistent with the explanation above

I think it's self-explanatory why everything is listed. The tree is a way to organize the bits of information you need for the problem in general. Not using things in the tree because they are the same (gravity, energy density of muscle, etc) comes later when you are putting everything together.

The muscle fraction and the energy density in muscle are also included, even though they eventually divide out.

These things are kept because they're needed to calculate an animal's jump height (which is what the tree claims to calculate). When calculating the _ratios_ of jump heights, these constants cancel.

The tree really should be put alongside or above the explanation. Obviously the explanation is needed, but a lot of times the assumptions are obvious, and tedious to read through

I find this tree really helpful...its easy to forget steps as you move down

$$E_{avail} \propto m_{muscle},$$

where this last step uses the assumption that all muscle has the same energy density $\mathcal{E}$.

Here is a tree that summarizes this model:

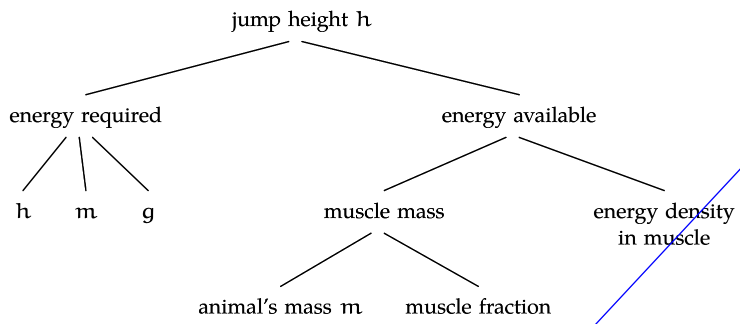

Now finish propagating toward the root. The available energy is

$$E_{avail} \propto m.$$

So an animal with three times the mass of another animal can store roughly three times the energy in its muscles, according to this simple model.

Now compare the available and required energies to find how the jump height as a function of mass. The available energy is

$$E_{avail} \propto m$$

and the required energy is

$$E_{required} \propto mh.$$

Equate these energies, which is an application of conservation of energy. Then $mh \propto m$ or

$$h \propto m^0.$$

In other words, all animals jump to the same height.

that seems intuitive, no?

this method does seem intuitive. i think that it is because it is stuff that we have learned before

for me, the method is not at all intuitive. you could argue the final result is intuitive, but personally I would not have come up with these steps for approximating on my own.

Still, we are only looking at the muscular power of an animal. Isn't it possible that animals that are larger have to have an exponentially higher amount of muscle to exert the same amount of power on their legs?

This is true but we are focusing on animals in general and if we can use various specific animals to get an estimate we can make an approximation for all animals.

Which would make sense without using equations and just common sense

$$E_{avail} \propto m_{muscle},$$

where this last step uses the assumption that all muscle has the same energy density $\mathcal{E}$.

Here is a tree that summarizes this model:

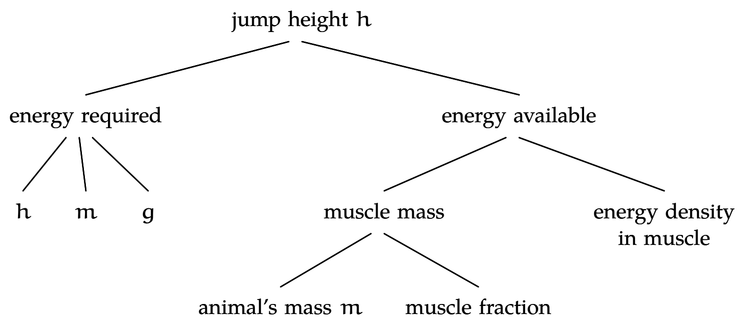

Now finish propagating toward the root. The available energy is

$$E_{avail} \propto m.$$

So an animal with three times the mass of another animal can store roughly three times the energy in its muscles, according to this simple model.

Now compare the available and required energies to find how the jump height as a function of mass. The available energy is

$$E_{avail} \propto m$$

and the required energy is

$$E_{required} \propto mh.$$

Equate these energies, which is an application of conservation of energy. Then $mh \propto m$ or

$$h \propto m^0.$$

In other words, all animals jump to the same height.

How come some humans can jump higher than others, assuming both are athletes?

I think the idea that all humans can jump the same high is a very broad generalization. One athlete may have a different muscle fraction and thus muscle mass, and/or a different energy density in his or her muscles, which changes the energy available. Hence, they are able to use that higher available energy to overcome the larger energy required for a higher jump.

I agree...but we made a fair amount of assumptions here, so I feel like we aren't really learning anything new.... only that we can "roughly" jump the same height.

I think the point is more about how high different species jump relative to each other.

and from this class's perspective, all humans do jump the same height. whether you can jump 1 meter or half of a meter, it's all the same within a factor of about 10^0.3.

This is generalized 'to hell'; of course all people don't jump the same, but it's a really nice approximation, and a very good example here. I think adding some more animals would be nice, but there's really no way of going into more detail without going into way more detail about muscle structure and densities and it's going to get messy. This was definitely one of the clearest examples so far, for me.

This explanation makes a lot of sense. Maybe it shouldn't be all animals jump to the same height but all animals with the same muscles jump to the same height.

This is cool.

is this really true? is there a way to validate our approximation

this is because we modeled all animals the same

which we know is false- what is the use of "proving" this?

which we know is false- what is the use of "proving" this?

$$E_{avail} \propto m_{muscle},$$

where this last step uses the assumption that all muscle has the same energy density $\mathcal{E}$.

Here is a tree that summarizes this model:

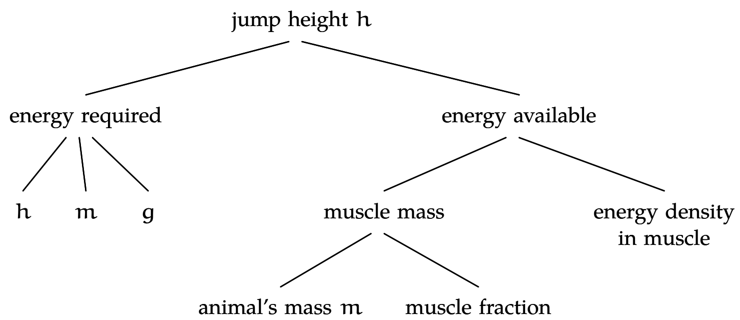

Now finish propagating toward the root. The available energy is

$$E_{avail} \propto m.$$

So an animal with three times the mass of another animal can store roughly three times the energy in its muscles, according to this simple model.

Now compare the available and required energies to find how the jump height as a function of mass. The available energy is

$$E_{avail} \propto m$$

and the required energy is

$$E_{required} \propto mh.$$

Equate these energies, which is an application of conservation of energy. Then $mh \propto m$ or

$$h \propto m^0.$$

In other words, all animals jump to the same height.

I wouldn't have guessed this.

Me neither, but I think some of our more questionable assumptions (for example, all animals are equal percentage muscle, which I find suspicious) probably brought this result about.

I agree, I believe this result is mostly a function of our assumptions.

I think this is just a function of the assumptions. Most importantly that energy and mass scale with each other. I might be interesting to estimate the amount of mass that an animals legs make up and incorporate that....

It's not supposed to be an exact approximation, but it does imply that a flea would not be able to jump 10 times higher than a human given the respective muscle masses. The approximation is rough but pretty accurate given our generalized assumptions.

this seems strange to me too. it also seems false.

This seems sketchy to me too but I think the idea is that, with the exception of a few very extrordinary animals, all animals jump about the same height.

Maybe all animals of equivalent mass can jump the same height, but all animals cannot jump the same height. That simply doesn't make sense

I feel as though this is an example in which there is too much error approximation, and the true answer (which is within a small number of magnitudes) has been lost.

I agree, this feels similar to the plane range problem in the previous reading. A lot of assumptions were made, and then it was just declared that all planes travel the same distance. Similarly, here, a lot of assumptions were made, and then it was just declared that all animals jump the same height. However, I think it would be better to remind the reader that this is the result for comparing 2 animals with similar mass, muscle percentage, etc. You're basically comparing 2 identical animals, so it's no wonder they can jump the same height.

this is amazing

I feel like this conclusion is analogous to the plane/bird one. We learned that birds can fly about the same distance as planes, and now that all animals can jump to the same height.

I'm a little surprised by this conclusion.

The result, that all animals jump to the same height, seems surprising. Our intuition tells us that people should be able to jump higher than locusts. The graph shows jump heights for animals of various sizes and shapes [source: *Scaling: Why Animal Size is So Important* [26, p. 178]. Here is the data:

Animal	*Mass* (g)	*Height* (cm)
Flea	$5 \cdot 10^{-4}$	20
Click beetle	$4 \cdot 10^{-2}$	30
Locust	3	59
Human	$7 \cdot 10^{4}$	60

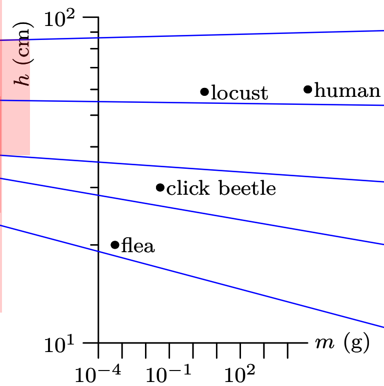

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

This is a pretty impressive result and a really cool example!

this is confusing- I think you mean total height, while most people would think this means proportionally

Really? I would guess that locusts can jump higher..

I like this follow-up with some numbers to check our calculations. This problem in particular does a good job of it.

I don't know about this. Locusts are pretty "jumpy" Maybe we'd expect to jump higher than, say ants, or things you don't normally see jumping anyways.

Is this graph intending on showing a general trend among all animals? For example, do elephants jump really high?

This makes sense with the assumptions we made but I feel as though it takes the mean of all animals when in reality different animals have been able to survive due to their ability to exercise the strengths they have, such as the fact that ants can hold much more weight then themselves whereas humans don't have the same ability.

This conclusion isn't trying to say that ants can jump as high as locusts. That's too specific of a comparison. Ants weren't designed to jump, so you're comparing apples to oranges.

What the conclusion is saying is that out of all the creatures that do take advantage of jumping, they jump to about the same height, regardless of how much they weigh.

It is a very broad estimate on a theoretical limit of how high living things can jump, given physiology of living things and the chemical power of muscle matter.

Actually, that's a really good point that makes all this a lot more believable to me. I think it'd be good to have something like that in the book, too?

The result, that all animals jump to the same height, seems surprising. Our intuition tells us that people should be able to jump higher than locusts. The graph shows jump heights for animals of various sizes and shapes [source: *Scaling: Why Animal Size is So Important* [26, p. 178]. Here is the data:

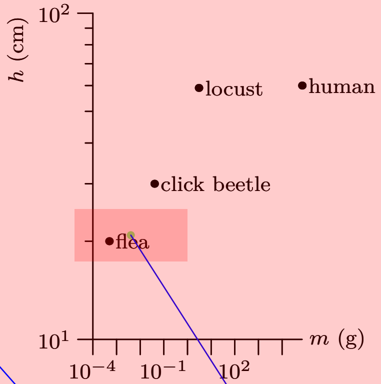

Animal	*Mass* (g)	*Height* (cm)
Flea	$5 \cdot 10^{-4}$	20
Click beetle	$4 \cdot 10^{-2}$	30
Locust	3	59
Human	$7 \cdot 10^{4}$	60

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

I think it might be good to add the kangaroo on this graph seeing that you specifically mentioned it before. For those who are interested, after a quick google search, a kangaroo can jump 10 ft or 300 cm.

Is that how high a kangaroo can jump from a standing position? The data in the table is for a "standing high jump". In a running high jump, you get to store energy in your motion and then convert that into height. In the standing high jump, all the energy has to be stored in your muscles (or shell, if you have one).

I understand the argument being made, but i have to disagree. Lets put it this way. Muscle fiber, regardless of the size of animal it is in, has the same strenght (as the cells that make it are always the same. The strength of muscle depends on cross-sectional area (call it l^2) but the weight depends on the volume (l^3). This implies that smaller animals (even ones with the same muscle-weight percentage) will be stronger per body weight than larger ones. this is why ants carry so many times their weight, and why elephants cant jump. THis should also imply that smaller animals should be able to jump higher in proportion to thier size than humans. IS this reasoning incorrect?

So why can ants carry so much more than a human??

Wait, how does that relate to jumping?

Hahaha it doesn't, I was just wondering if similar principles apply?

It probably does.

given how they plagued my childhood pets, i'm a bit underwhelmed by how they stack up.

This last section was really helpful in understanding the previous part. I'm really glad that this section is included.

I'm not sure about the accuracy of this. I think fleas have about a 6 foot max range - so about 180 cm.

That's horizontally, though. This is vertically. I'm not sure how that affects your number.

I like that you make assumptions generally in the beginning and then work through the calculations, then finally provide data.

The result, that all animals jump to the same height, seems surprising. Our intuition tells us that people should be able to jump higher than locusts. The graph shows jump heights for animals of various sizes and shapes [source: *Scaling: Why Animal Size is So Important* [26, p. 178]. Here is the data:

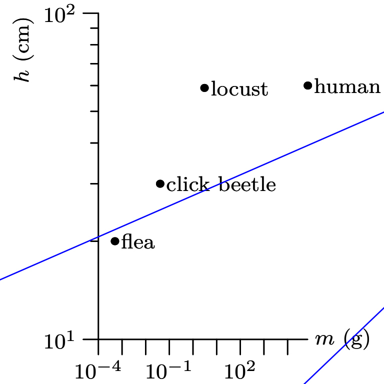

Animal	*Mass* (g)	*Height* (cm)
Flea	$5 \cdot 10^{-4}$	20
Click beetle	$4 \cdot 10^{-2}$	30
Locust	3	59
Human	$7 \cdot 10^4$	60

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

I think penguins should be here, I read somewhere that they are unique in jumping and it would be nice to analyze them

wait–are you saying that jump height is proportional to mass, or that the variation in the data is significant? i'm confused.

At first I thought it was saying that height is proportional to mass but then it goes on to say that jump height variations are much smaller than the variations in mass, so I'm not really sure.

Right, it just means that the height doesn't change nearly as much with huge changes in mass (i.e. a large animal still jumps about the same height as a small one). I like this observation, it acknowledges the rough approximation but still reveals an interesting fact.

The result, that all animals jump to the same height, seems surprising. Our intuition tells us that people should be able to jump higher than locusts. The graph shows jump heights for animals of various sizes and shapes [source: *Scaling: Why Animal Size is So Important* [26, p. 178]. Here is the data:

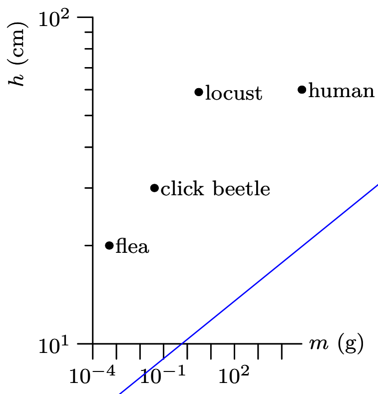

Animal	*Mass* (g)	*Height* (cm)
Flea	$5 \cdot 10^{-4}$	20
Click beetle	$4 \cdot 10^{-2}$	30
Locust	3	59
Human	$7 \cdot 10^{4}$	60

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

What about animals that can't jump? I don't think, for example, turtles or elephants or some things even can jump...

Or starfish.

It looks like everything listed there is also an insect - if we throw in things like lions and kangaroos, wouldn't that throw things off?

Not sure, but I agree it would be nice to have data from something else besides people and insects...

I think that's one of the problems with the simple model.

I thought we were going to make some corrections to this simple estimation that would refine it to handle other mammals?

I don't think turtles or elephants jump at all, so maybe we're restricting our analysis to animals that jump regularly or have the ability to jump regularly?

well i mean, if we say turtles can't jump at all–well no fish can jump so obviously we're not talking about fish. but yeah, if we got into the case oflarger animals like elephants or giraffes or something, we might have to make a new category–maybe different categories of jumping height would make the result a little more realistic

fish do jump...

read the above thread for an answer to this thread

This may have something to do with the simple spring model of jumping we assumed at the beginning of the section. It's probably only reasonable to apply it to animals that can more or less jump up and down.

Is it that elephants and turtles physically can't jump, or they just...don't?

The muscle fraction issue comes into play with things like elephants, I think. Bone strength is proportional to cross sectional area (L^2), while mass goes as L^3, so heavier animals need comparatively larger bones to support their own weights. I suspect that you would find that the muscle fraction of an elephant is lower than that of a kangaroo.

ok, it answers my question here =)

The result, that all animals jump to the same height, seems surprising. Our intuition tells us that people should be able to jump higher than locusts. The graph shows jump heights for animals of various sizes and shapes [source: *Scaling: Why Animal Size is So Important* [26, p. 178]. Here is the data:

Animal	*Mass* (g)	*Height* (cm)
Flea	$5 \cdot 10^{-4}$	20
Click beetle	$4 \cdot 10^{-2}$	30
Locust	3	59
Human	$7 \cdot 10^{4}$	60

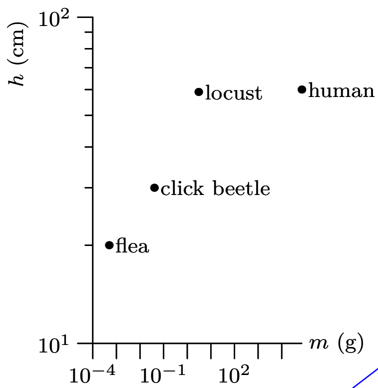

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

While this does make sense, I feel like this contradicts the previous statement saying all animals jump the same height. I understand that they all jump with similar heights, I just don't like how the example went out of its way to "prove" that all animals jump the same height.

So i guess this is where the whole issue of how we can compare such differently shaped and sized animals together. I was wondering how we take into account the clear advantages larger animals have. But spread out over a graph, our guess was right.

this seems really cool, but I am having trouble using this method in my own calculations- how to I know how to take it- should I start with a tree or start by just figuring out proportions

Sometimes, the first time I read these kinds of parentheticals, I think they are products...

Is this the first time we've seen this notation? I think it's confusing since 1 is unitless, though I realize h is proportional to mˆ0 and things to the zeroth power are 1.

Sometimes, the first time I read these kinds of parentheticals, I think they are products...

Sometimes, the first time I read these kinds of parentheticals, I think they are products...

I thought this was a great summary.

So basically, per unit mass, all animals can jump to the same height?

I think it'd be cool to bring back the gravitational constant g, just to prove how height varies with it. It'd be cool to prove exactly how high people on the moon can jump using proportionality between the earth and moons gravity. Just an idea...

I think the problem most people have with this is that yes - the order of magnitude for the jumps is all the same relative to mass, but when we consider jumping, a factor of 2 seems like a LOT to witness physically, but not all that much in this calculation. So while yes, we all jump the same magnitude of height, there are some differences...is this correct?

Hold on though, the jumps changed by a factor of 3, while the masses changed by a factor of 10ˆ8. To me that makes me believe our calculations.

The result, that all animals jump to the same height, seems surprising. Our intuition tells us that people should be able to jump higher than locusts. The graph shows jump heights for animals of various sizes and shapes [source: *Scaling: Why Animal Size is So Important* [26, p. 178]. Here is the data:

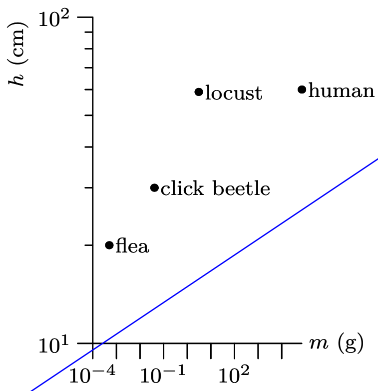

Animal	*Mass* (g)	*Height* (cm)
Flea	$5 \cdot 10^{-4}$	20
Click beetle	$4 \cdot 10^{-2}$	30
Locust	3	59
Human	$7 \cdot 10^4$	60

The height varies almost not at all when compared to variation in mass, so our result is roughly correct! The mass varies more than eight orders of magnitude (a factor of 10^8), yet the jump height varies only by a factor of 3. The predicted scaling of constant h ($h \propto 1$) is surprisingly accurate.

it'd also be cool to see what the limits to our jumping capability based on the average weight and range in weight a human has

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

Interesting...is this primarily energy from glucose or carbohydrates or possible energy that can be stored in fats? I can imagine energy from fats will take longer to break down.

The homework is very useful in understanding the 2nd problem.

Is this the animal's size in height only or its size in weight, height etc?

I think size here usually refers to the length dimension. Weight scales as the cube of that, since W mg volume L^3. It's that difference in scaling that provides the interesting findings since things scale differently relative to L and m.

It is so weird that small animals need more power and have higher drag. It seems quite the opposite, when you compare the speed of a fly to say a human.

They don't need higher absolute power, just higher power/mass ratio. And they don't have higher absolute drag, just higher drag energy/kinetic energy.

Is this because the small animals need by energy? Could you explain in more detail how you made these connections?

Is this true for all animals? Or specific for fleas? The height proportion is confusing me...seems wrong to estimate jump height as 1m.

I think this was a very clear and descriptive memo that made it easy to understand. I am starting to feel more comfortable understanding the calculations and examples that are given.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

So, I am still lost with this whole proportionality thing; can we get a table of proportionality, or a breakdown of exactly what goes into it? I'm still stuck thinking in units, and that's not helping me at all. I see t prop length, and m proportional to l^2. Is everything proportional to some sort of l?

Not everything is proportional to L, for instance, we found that jumping height isn't, since the L's cancel. You're right that units can get confusing if you try to think of all of them at once. Force is proportional to mass, but it's also proportional to acceleration. Acceleration is proportional to distance, and also inversely to time squared. Instead of including all the units at once (kg-m/s^2), you can think of them affecting force independently (mass^1, distance^1, time^-2) or in whatever combination you want.

This section reads very well. The transitions between subsections and topics helps me to anticipate the information I'm about to read and helps me to understand it when I read it.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

Comments on page 1

Read about power limits and fleas (two subsections) for the Thursday memo.

I feel like this section just jumps into the math much more quickly than the other sections. In addition, I feel like the math is not being explained as much.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

I thought about this in the last reading but didn't comment about it. Athletes with strong legs are not the ones who jump high, athletes who are "explosive" do. In sports, this means being able to provide a lot of force in a short amount of time, which in physics is obviously power. To summarize, I expect power to be a much better estimator of jump height than energy.

Slow twitch versus fast twitch

great way of saying this!

I don't know about slow twitch versus fast twitch... Baseball is the best example of a fast twitch sport, fielders stand around for minutes between sprints and hitters do the same. I would guess they have the worst verticals of all athletes.

That's because baseball doesn't require you to jump, but I bet they all have a bank robbers first step.

Also, I would strongly disagree with the verticals comment. Think about, for instance, Derek Jeter, or some similarly athletic player. Infielders especially need to jump high for line drives, something greatly effected by vertical jump.

I think that's debateable per se. For instance, slow twich muscles could still be stronger than fast twitch in two relatively similar individuals depending on training and composition. For instance, sprinters often do weight lifting on legs for training to gain speed, but I know some of us long distance runners are actually much faster out of the blocks.

That's unusual, in my experience. I'd say a well trained sprinter should beat out a well trained long distance runner any day. Although, do you mean faster to leave the blocks or going faster upon leaving them (or shortly thereafter)? The first is largely a function of reaction time, whereas the second is more a function of fast vs slow twitch and muscle strength.

Slow twitch muscle fibers aren't "stronger" than fast twitch muscle fibers and vice versa. The slow/fast twitch simply means how quickly individual muscle fibers are recruited to contract when a signal from the brain is given. Having more fast twitch results in recruiting more fibers at once (power) as opposed to slow twitch which recruit fibers over a longer period of time.

I think it was said earlier but these fiber affect power and not strength. You can lift weights and strengthen both muscle fibers....

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

I think this answers my question from last time about why some people can jump higher than others

This is probably why personal trainers recommend rapid repetitions over slower ones.

It depends on what you're doing, fast twitch muscle training involves quick repetitions which we need for explosive motions but if you are trying to build strength then heavy slow repetition is desired. For this example fast twitch muscles will get you more height but if you are trying to squat then you wouldn't want to train the same way.

This is something I never expected to learn in an estimation class!

I'm unsure of what you would not have expected. The reason it has to be rapid is to increase the power. Did you not know how to calculate power?

Kinda random but related to delivery speed: Read an interesting article somewhere about sprinters tending to have shorter moment arms for the achilles tendon. Though this decreased the moment that could be applied, it extended the duration of the "flexxing." Apparently muscles are better at applying force over time and this outweighed the negatives, from a reduced moment arm.

I think it would be nice if you wrote out "jump height limit" here to clarify what limit you are talking about.

yeah i wasn't sure what "limit" we are talking about here. Energy limits a lot of things that we can do.

Did you mean to have both words here? Could you just pick one?

I'd prefer 'varies'

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

Is the variable, l, the length/size of the muscle or animal?

I assume animal simply based on previous examples, but it isn't very clear.

Yeah it's not explicitly said here. It could be muscle size or animal height, and should be defined.

How big a difference would it make?

we've always used it as a characteristic length (or length scale) of the animal.

I also assumed it was size of the animal, because the previous paragraph doesn't mention anything on the size of the actual muscles, only how much energy the muscles can store.

@5:05 ... it would to the actual approximation, just to the general understanding of the process

just missing a "."

Where? I'm not sure what you're talking about.

between the "l" and the capitolized "Power"

I agree.

What about if the animal is designed to maximize the distance of its jump, can we consider that as well? I feel like maybe a grasshopper or something would jump far instead of high.

I was wondering about that as well... Is there any correlation between how high animals can jump and how far they can jump? If we assume everything is a cube, then that factor doesn't change in any direction.

I'd imagine there's a correlation between jump height and standing broad jump. However, if you take into account a run up like in the long jump, for example, speed becomes a major factor.

I'm sure there's a correlation. Since jumping forward and high involve a lot of the same muscles, it would be intuitive that they go together.

I'm sure there's a correlation. Since jumping forward and high involve a lot of the same muscles, it would be intuitive that they go together.

Is this the time before or during the jump?

During, i.e. while the muscles are delivering the power to propel the animal upwards.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

Trying to apply this proportion to humans... m (2m)^3 seems a lot bigger than a human's actual volume...

True, but maybe it gets close enough since we would just round it to few or 10^1 anyway...

Why is m l^3? Shouldn't we account for density? I'm assuming that p = 1 which is why m l^3, but could someone please explain this to me?

Density for all living creatures is usually bout water weight

...and volume is dependent on meters squared.

Er, no, volume has units of meters cubed. Density is the reason why it's only proportional to, and not equal to. We're not really assuming rho is 1, we're just ignoring it entirely, assuming it's pretty much the same for all animals.

Thanks, that makes sense now: rho is ignored since it's the same for all animals.

In that case, it might be helpful if that was explicitly stated in the notes. I know it may seem like overkill, but I would have thought it beneficial.

To clarify the above comment, m is proportional to l^3 partly because we are omitting a constant density, but also because a person or animal's volume isn't necessarily equal to l^3. l is just a characteristic length that we can use to compare the sizes of different animals.

it's talking about proportionality. mass is proportional with size. density is an unchanging factor.

Shouldn't we be calculating "l" as say the length of a person bent and folded into a cubical shape, like he did in the class example?

Yeah, that's what I would guess. I think we used the value of .5m in class to represent l for a human. It probably evens out the same as if we were to calculate l*w*h but in a more simple way.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

why a squiggle? isn't this a law of physics?

I was wondering that too. But if I had to guess, it would be because the practice of estimation has so ingrained in us a mentality to go for proportions and relations rather than exact equations, that even when faced with a known equation, the habit of using squiggles still carries over.

I thought that just means we don't have to worry about any constants.

There are a lot of factors and terms that we are neglecting. Drag, muscle efficiency, energy spent on balancing, etc.. But roughly speaking, E goes like mgh.

I get that we don't have to worry about them when we are trying to find out proportionality, but if we later use the results to plug in numbers, do we add them back?

This section is about proportionality, if you haven't noticed, all physics related equations are simplified to a proportionality.

how is h independent of l, I thought l was length I feel like they would be connected

It would be really helpful to have a table of all the proportionalities. I keep forgetting the ones from the previous sections.

I feel like having a large table of proportionalities would kind of defeat the purpose. The ideas is to gain experience with proportional reasoning so that we can apply it to many situations, not just those situations which have already been covered.

Well if the table of proportionalities would kill the purpose, can we at least get a table of equations?

I agree with not having a table. I feel like once you take a moment to really think about it, something clicks and you understand it. Rho is same for all animals so the relationship exists in this case. However I find that I need to see it once - it's not natural yet but practice will only help so I wont to be forced to think about it.

how would you even estimate this time for each type of animal?

even when it's straightening its legs and pushing against the ground itself?

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

A previous comment suggested that the time over which the energy is delivered occurs during the jump, not before. However, this makes it sound that the energy comes from before the jump until takeoff, which is the case?

Are we assuming that the same force is being exerted on the animal whenever the animal is on the ground? It would seem to me that this force is related to the surface area of the animal that is touching the ground. So, as the animal gets closer to leaving the ground and only it's toes, for instance, are touching the ground, less force is exerted on the animal.

I'm pretty sure that we can assume that the force applied is constant throughout the whole jumping process.

This is really interesting. I have always wonder how you can do measurements like these.

From 2.671 go forth, the constant force approximation for jumping is not far off from what's actually happening. For those curious, this is pretty easy to measure using a force plate.

So is delivery time the same thing as contact time here?

Yes, I believe that is what this is saying.

Would it be better to say distance here since that's the more typical variable and the following equation uses "extension distance" anyways?

i don't really understand this value. what are we defining as "extension" here. The crouching down and then back up again?

Yes. I believe the "extension length" is the difference between crouched height and fully extended height, which is the distance over which the animal exerts force on the ground.

This explains my earlier question about delivery time.

I like this method of finding the "lag time" between telling your muscles to jump and actually leaving the ground.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

How is the extension distance roughly the animal's size? If extension is simply motion before leaving the ground, it would have to be more like half the animals size.

I agree, maybe it's since we're just estimating and we know it's not going to extend any MORE than its own size.

Movement through the distance is not constant, but goes through a lot of acceleration. Are we trying to take an average speed here?

Yes, I believe we are taking an average speed just to simplify the calculation.

So we are allowed to make this approximation because the extension distance is proportional to the animal's size? Therefore, we can throw out any constant factor like extension distance is 1/2*L?

I feel like this is a little funny because there is certainly a constant factor being ignored here. In the end, should be same order of magnitude...

Ahh, the magic of the "proportional to" sign.

Yeah, in the last few readings we've ignored the constant terms. this is because we're interested in proportionality and not concrete values.

I agree with you guys, it reads a little funny but still makes sense. Maybe it would be a bit clearer if it notes the extension distance is a function of the animals length, and then use that proportionality?

It seems to make sense: think of a cat jumping from a crouch. Right before they leave the ground they're fully extended with only their back legs on the ground, so they're at their full height. Then the extension hieght is that full height minus their crouched height... and you can assume that more most animals their crouched height is some fraction of their full height and assume that that fraction is the same for most animals... So the proportionality works out. I think.

But how does that work for humans? I could believe that extension length is maybe half of our body size. This is counter intuitive if you picture someone jumping. Ah approximation.

However, if you consider the l to be 0.5 m like mentioned earlier when we were talking about human volume then it is not that difficult to imagine extending you legs from a folded stance about 0.5m.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{delivery} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{takeoff} \propto h^{1/2} \propto l^0.$$

So

$$t_{delivery} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

I wouldn't have thought this was so similar at first, but when I think about it now - it does make some sense (to first order of course), but I'm trying to think about what factors go into this.

I would have thought this was a bit faster than the extension speed, just because I feel animals undergo an acceleration while extending, and not quite a constant velocity.

Yeah, but you push off the ground when you jump. Are we assuming that that pushing contributes to the acceleration in the extension or the jump itself?

Animals do undergo an acceleration when extending, but my thought is that it the acceleration happens in very little distance, so the average speed might be comparable to the takeoff velocity. What do you think?

What model is this referring to?

why can we make this assumption? it seems like takeoff velocity would vary a lot among different animals

Yeah I also feel like this could affect the delivery time.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

Where does this come from? Is this from an earlier reading? If not, I think it needs more explanation because I'm confused as to why this is.

I agree. I don't think it is from a previous reading. And its not a clearly intuitive relation that we can just assume.

A mention of the conservation chapter and kinetic/potential energy conservation would be nice here, since it's a concept that comes up a lot, like in the roller coaster design example.

I too am a bit confused where this came from exactly

I think this was mentioned briefly before, but a refresher would be nice. I believe it comes from setting the kinetic takeoff energy to the potential energy of the jump height: v^2=gh. That v_takeoff is the same for all animals follows from our conclusion that all animals jump to the same height.

can you conclude that the takeoff velocities are the same from the fact that they jump to the same height?

Generally, I would say yes. If we set KE=PE, mv^2=mgh, cancel the m's and v is proportional to h^1/2. This is the equation he has here, and this is where it comes from. It makes sense since all animals have the same acceleration to contend with, they must start at similar velocities to get to similar heights.

great explanation, perhaps a little refresher here would still be good though.

great explanation, perhaps a little refresher here would still be good though.

great explanation, perhaps a little refresher here would still be good though.

great explanation, perhaps a little refresher here would still be good though.

great explanation, perhaps a little refresher here would still be good though.

great explanation, perhaps a little refresher here would still be good though.

Great comment. This helped me understand.

confusing derivation

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

I'm not sure I quite understand the 0 exponent...am I missing something obvious?

Me too... Doesn't l^0 = 1? Also what is h? The length extension? I fee like the variables here weren't defined properly.

The 0 is hard to read in this font at the magnification I normally use.

Shouldn't you just magnify it more then?

This is only a valid complaint, if when you print the file on actual paper (where you don't have the luxury of magnifying), it is still too small to read.

What is h again?

ohh yeah, it's the height

I like how you take time to spell this out. Walking through it really helps.

I understand that this comes from the above equation - it might be helpful to make this more explicit (as in, directly plugging in the proportions) to make this easier to follow.

It took me a few minutes to see that relationship. I agree that a quick step-by-step showing the plug-in values would help.

I think it definitely makes sense intuitively, smaller animals will take less time to extend in comparison to larger animals.

It's kind of ironic thinking about how something like power is only dependent on length. This method is pretty cool.

4.3.2 Power limits

Power production might also limit the jump height. In the preceding analysis, energy is the limiting reagent: The jump height is determined by the energy that an animal can store in its muscles. However, even if the animal can store enough energy to reach that height, the muscles might not be able to deliver the energy rapidly enough. This section presents a simple model for the limit due to limited power generation.

Once again we'd like to find out how power P scales (varies) with the size l Power is energy per time, so the power required to jump to a height h is

$$P \sim \frac{\text{energy required to jump to height } h}{\text{time over which the energy is delivered}}.$$

The energy required is $E \sim mgh$. The mass is $m \propto l^3$. The gravitational acceleration is independent of l. And, in the energy-limited model, the height h is independent of l. Therefore $E \propto l^3$.

The delivery time is how long the animal is in contact with the ground, because only during contact can the ground exert a force on the animal. So, the animal crouches, extends upward, and finally leaves the ground. The contact time is the time during which the animal extends upward. Time is length over speed, so

$$t_{\text{delivery}} \sim \frac{\text{extension distance}}{\text{extension speed}}.$$

The extension distance is roughly the animal's size l. The extension speed is roughly the takeoff velocity. In the energy-limited model, the takeoff velocity is the same for all animals:

$$v_{\text{takeoff}} \propto h^{1/2} \propto l^0.$$

So

$$t_{\text{delivery}} \propto l.$$

The power required is $P \propto l^3/l = l^2$.

That proportionality is for the power itself, but a more interesting scaling is for the specific power: the power per mass. It is

I am still not comfortable with being able to drop most of the values in an equation and say that energy required to jump is proportional to l^3 so let's only use this value. Why can you do this?

Proportionality is about the relations between quantities. We're not pretending that the power IS length squared, we're saying that as length changes, the power increases by the square of that change. One simple example I like to come back to is surface area to volume ratio, which goes as 1/L. Try to think about how the quantities affect each other, as if you were moving a slider for L and watching the change in the power. Ignore the absolute values, which are affected by constant, but constants are just that, they stay the same regardless of scale.

I don't know if you can include some sort of comments section in the book but I have found that these comments really make the material much more transparent and easier to understand. I think finding a way to compile/summarize the comments and present them along with the book text would make the book GREAT!!! Maybe include them as little hints in the margins similar to the way some high school science books have random bits of information thrown in.

I think I'm beginning to really understand this idea now...to try to figure out proportionally how things compare, and commonly seeing how "per mass" factors in is quite helpful!

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

I think this was mentioned before too... having the equation and the introduction sentence on the same page is always better.

What is specific power?

A "specific" version of a quantity is that quantity divided by the mass.

I still don't get it.

I think it's basically the power it needs divided by its mass.

specific power= power/mass

specific power = Power/Mass. Just like specific heat = (Heat Capacity)/Mass, specific volume = Volume/Mass (inverse of density), etc.

This explanation helped out a lot...as someone that hasn't dealt with any of those terms in a long time the idea of specific power confused me

Thanks for the explanation! I think it would be helpful to put a single line in for this in the reading.

By dividing by mass, it basically "normalizes" the quantity to relative to the nature of the object.

is specific power inittial power?

maybe its same as power density or power per unit of mass

yeah, what he's saying here is that we should look at specific power–power per mass. this happens to be proportional to 1/l since p is proporitonal to area (or l^2) and m is proportional to volume (l^3). divide area by volume and you get 1/l.

This sentence is just confusing to me.

Since P/m is proportional to L^-1 it means that smaller animals (those with small "L's") will have higher P/m to generate the P required to jump. Hope this helps.

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

This is somewhat surprising.

Really? While reading this section I had a mental image of a lever arm in contact with the ground. The longer the arm, the more contact time, so the less specific power needed.

I also remember hearing from elementary school or maybe Bill Nye that if we had the same power as an ant proportional to its size, we could lift cars above our heads.

i <3 bill nye

I don't know how the Bill Nye comment relates to this context.

This passage is saying that smaller animals need a higher specific power to reach the same height, putting them at a disadvantage. The Bill Nye comment says that smaller animals, like ants, have it easier.

I think Bill Nye was referring to energy. If that were the case, then he'd be right. In fact, we would probably be able to lift more than a car, which is only the weight of 10 times our current lifting capacity.

Agreed, it makes sense, but you have to think about it more

What exactly is considered a "small animal"?

I think it's just setting up a relationship (i.e. as size decreases, required specific power increases.)

We've talked about animals as small as fleas, so I guess this could refer to bugs?

Does this ratio also show that smaller animals are more capable of moving quicker and generally being stronger because mass will have less of an affect on power because of the dimension ratio?

It does seem to be correlated. It is interesting though, that you can consider animals like ants that can carry items many times their weight.

Or as we see below, jumping fleas!

at what point does a "smaller" animal do have enough power? for example, are frogs big enough?

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

do you mean large enough?

i agree that this is confusing. i think that it's saying that if small animals need more power, but we all produce the same amount of power, so they wouldn't have enough. it depends on your reference point i suppose.

he's not arguing at all that all animals produce the same amount of power (and that would be false). Power required scales like l^2, but specific power required scales like l^-1. So as l decreases, specific power required increases.

Small here refers to length, not mass. (Yes, I know these things have their own relationship, but the key is that P/m gets too high to achieve physically for animals as their L decreases.

this still doesn't seem right- large animals are not always the highest jumpers

I guess speed is included in the analysis, but it seems as though time specifically should have been added, since it's analyzing the amount of time over which energy is released.

This is interesting - why does this apply only to small animals?

Why is it then that some bugs can jump so high?

Never mind I should have read a couple sentences further

why is this? does is relate to muscle size?

I'm not sure if this is correct but maybe we can make an analogy to gears. If a small gear wants to release the same amount of energy as a big gear it has to spin really fast. Small animals, since they have smaller "L's" can "spin fast enough" to keep up with animals with bigger "L's". Therefore they can't release their energy fast enough and thus can't jump as high.

Did I miss this graph and data set?

Its at the end of reading 15.

I believe it was in the previous reading memo.

Sort of like skiing, you get energy out of the camber of the ski and then it releases the energy to shoot you into the next turn.

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

that's really cool, but just goes to show that our estimations don't account for everything, and is not a reliable.

Rachets maintain posistion and/or prevent movement in the opposite direction, but not necessarily with bending or stored energy. This is more like a leaf spring.

This seems like it should either be made more important and removed from parentheses/explained more thoroughly or dropped from the text - as-is it doesn't seem to add to the example and adds one more thing to think about

Doesn't that depend on how one holds it motionless? Letting a weight sit on a table requires no energy, but holding it above the table does?

To me, a ratchet mechanism indicates not requiring energy to hold something motionless. Isn't that why people often use ratchets? To keep a shaft or a rope or something else from slipping without requiring energy?

Unless you're making a subtle distinction between energy and power (in some cases takes initial energy but no constant power to hold something motionless), I don't see how this statement is necessarily true.

Could "energy" be referring to potential energy? in the tension of the shell or whatever restoring force is acting on the beetle's shell when it is ratcheted out of resting position. That would make more sense to me.

interesting way of relating this example to a well-known mechanical tool!

what is specific power limit again?

That's what he just derived: P/m = l^-1

along these lines, I always wonder where are the bounds for which these "animals" actually follow the predicted outcome (i.e. jump height etc.) all I can think about are the exceptions, I would like more examples of what animals they actually work for

so I looked these up on you-tube and they are pretty awesome

thats amazing

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

That's really clever.

This is cool to think about. How much difference does this alternative energy source make in the beetle's overall jump height?

Wow, that's really awesome!

It's an alternate power delivery system, energy is still conserved, it's just storing it up over a longer time and releasing it over a shorter one, like a capacitor, to bring in the electrical analogy from earlier in the course.

so then I guess since this release in energy is large compared to the total energy needed to propel the beetle far (in relation to its body size), it can jump really far. so that's why larger animals can't use this right–the energy supplied from this type of shell snapping is actually small, but large compared to the energy needed to move the beetle

there are a lot of really sweet videos of the beetle...one of the first ones I saw is http://www.youtube.co

Nice find! Interesting solution to the energy limit.

awesome. this sounds really interesting.

Awesome. I was just going to go look for something like this.

This is really cool! Are there other animal examples of this phenomenon?

Would the number of legs be proportional to the amount of power that can be delivered? For example, would increasing the number of legs allow for greater power delivery?

It might be nice to say that small animals can't jump as high rather than just calling it a discrepancy.

It's a discrepancy between our calculations and actual observations.

I like that the concepts we learn earlier come back and are linked to the current topics.

yeah I agree, it's a really good way to tie all the concepts together and make sure we don't forget them!

I feel like this comes up really often in our examples. Is that just this unit, or approximation in general or this type of problem or what?

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

This always seems to come back!

And i never would have considered it in this case, were it not presented to me. Although when I think about it, it makes sense that light enough animals would be effected.

Yeah this also seemed to come out of nowhere for me too. You figure we would get better at knowing when to consider it since we've used it a few times already.

So since the area scales with L^2 are we going to show that drag is the limiting factor for larger animals? Wouldn't the surface area to volume ratio go up for small animals and then they would have a larger effect from drag?

Huh, that's surprising, I would have thought it to be the opposite, since smaller creatures have less surface area...

Yes, but I think they're also less dense. (Though, come to think of it, if we're all comprised of the same materials, then it doesn't matter). But something tells me the very very small mass, relative to the surface area causes a problem. F pAv^2, but p m/V, so m/some L (thickness maybe?) probably shows that m increases much faster than an L. Just my two cents; I could be completely wrong.

I would have thought the exact same thing.

but its interesting that the force (in the equation above) has no dependence on mass. or do you mean small in terms of size? then it makes more sense because F depends on the area

I think this refers to area.

don't small animals have less area, so why then does the drag force affect them more

My first reaction here was: "but why", and then thought about it for a minute or two before realizing you explained it below. Of course that's my fault, but it gave me pause while reading it.

For the first time ever, we were given a bried synopsis or abstract of the theory about to be presented. I like it.

$$\frac{P}{m} \propto \frac{l^2}{l^3} = l^{-1}.$$

Ah, smaller animals need a higher specific power!

A model for power limits is that all muscle can generate the same maximum power density (has the same maximum specific power). So a small-enough animal cannot jump to its energy-limited height. The animal can store enough energy in its muscles, but cannot release it quickly enough.

More precisely, it cannot do so unless it finds an alternative method for releasing the energy. The click beetle, which is toward the small end in the preceding graph and data set, uses the following solution. It stores energy in its shell by bending the shell, and maintains the bending like a ratchet would (holding a structure motionless does require energy). This storage can happen slowly enough to avoid the specific-power limit, but when the beetle releases the shell and the shell snaps back to its resting position, the energy is released quickly enough for the beetle to rise to its energy-limited height.

But that height is less than the height for locusts and humans. Indeed, the largest deviations from the constant-height result happen at the low-mass end, for fleas and click beetles. To explain that discrepancy, the model needs to take into account another physical effect: drag.

4.4 Drag

4.4.1 Jumping fleas

The drag force

$$F \sim \rho A v^2$$

affects the jumps of small animals more than it affects the jumps of people. A comparison of the energy required for the jump with the energy consumed by drag explains why.

The energy that the animal requires to jump to a height h is mgh, if we use the gravitational potential energy at the top of the jump; or it is $\sim mv^2$, if we use the kinetic energy at takeoff. The energy consumed by drag is

This part reads a bit awkwardly. Perhaps remove the commas before the 'if's. Or, in the spirit of course 6, write it as if...then... statements?

Or break it out of sentence form into a set of two equations with notes: E mgh using PE at peak E mv^2 using KE at takeoff That would keep it in line with the next part about Edrag.

I like the second solution here. I think as long as this is in sentence form then it'll be awkward. It's a lot easier to follow if it takes the form of the Edrag equation.

This is a good point. I'm not sure I would have made this connection left to my own devices.

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\mathrm{kg\,m^{-3}}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\mathrm{kg\,m^{-3}}$. The

Still unsure about why the h is here.

This equation is looking at the energy lost from the drag force not the the force itself. This seems pretty simple to me.

Because energy is force through a distance (E = F * d). Since we are trying to find the energy consumed by drag, we use the drag force: E_drag = F_drag * d.

keep in mind that the drag force is changing with height, and that it also happens to be the case that integral(F*dy) v^2*h, as well.

As a course 6 major, I'm unfamiliar with a lot of these calculations about drag. I was curious if drag is affected by gravity? For example, would a jumping animal experience a different amount of drag on the moon? Is it that all animals would still jump to the same height on the moon, but that height would be higher on the moon than it is on the earth?

I think gravity figures into the density equation, so it is implicit that air density on earth is higher than say the ambient density on the moon, due to gravity.

Drag is a phenomenon that depends on an atmosphere. The moon has no atmosphere, therefore no drag exists. Think of it as friction from rubbing on gas.

Another way of thinking about drag is what Sanjoy mentioned one or two readings ago: The drag force is proportional to the amount of fluid displaced by the object.

That way, denser medium will yield a larger drag force, simply because you have to displace more mass per volume as you move through it.

thank you 10:33 & 10:52!

Agreed, thanks. This stuff is far less familiar to CS majors than the MechEs.

Total reversal from the beginning when we were talking about UNIX...

and those of us not meche or cs are still totally floundering.

That's a really cool way of measuring drag.

(and by measuring I mean thinking of.)

wow, this is great, I never thought about it this way.

wow, this is great, I never thought about it this way.

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\mathrm{kg\,m^{-3}}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\mathrm{kg\,m^{-3}}$. The

wow, this is great, I never thought about this way.

This type of ratio makes up a lot of important named parameters in fluid mechanics (Reynolds number, Prandtl number, etc.) Reynolds is kinetic forces vs. viscous forces. Maybe we'll come back to this in dimensional analysis.

wow, this is great, I never thought about this way.

Would you need some sort of coefficient in order to take the shape of the animal into account?

I think that trying the accuracy gained by trying to use a coefficient to make this animal shape more accurate than a simple box would be canceled out by all of the other simplifying assumptions we have made, and ultimately would not be worth the effort.

Unless the animal has holes in it or very different cross-sectional areas, pAh will still be the same mass of air swept out by its passing through. You're right, further refinement would involve things like birds and whales being streamlined, since that helps them not have to displace all the air they would if they were all cylinders.

Hmm this is also interesting and makes complete sense. Never would have thought of it that way.

Also gives a calculation why exactly jumping higher is harder....sort of

I really like this explanation of drag.

Yeah, this is really interesting...never would have thought of this...

right– it's a balance of mass to surface area.

yeah this sentence makes the idea of drag in this example really clear and easy to visualize.

This relationship is true for all size animals though

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\text{kg}\,\text{m}^{-3}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\text{kg}\,\text{m}^{-3}$. The

I don't understand how this assumption makes sense, unless you simply didn't want to have l^3/l^3

it makes sense based on the analysis we did previous (forgot whether it was in class or in the readings). it's not what we "want" it makes sense because the forces that we overcome by jumping are directly proportional to the capacity with we have to jump with...sort of

Just think about it in terms of a cube. Imagine them as little boxes with tiny legs. It's pretty standard physics to think of everything as a cube or a sphere.

I think this assumptions is too leading. Obviously, a beetle can't jump 2 feet in the air like a human can, but does that mean it doesn't have a proportional amount of power as a human?

What he's basically saying here is that the jump height is independent of the animal's "length" l. It might help if he explicitly stated it though.

I thought we were ignoring the density?

I feel like this will be dealt with once we start doing proportions in that it will probably be canceled out or be unnecessary.

we still are. the l comes from h

Is this supposed to be L^2 here?

I think so because right below, its L^2 on the numerator

I never would have guessed this, I would have guessed the more surface area, the more drag

You're sort of right; more surface area contributes more to drag, but heavier objects are intuitively more "immune" to drag. So it really comes down to surface area to volume ratio. Smaller animals have a larger area to volume ratio, despite having a smaller area overall.

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\mathrm{kg\,m^{-3}}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\mathrm{kg\,m^{-3}}$. The

If this is true, why can smaller animals often jump so much farther than larger ones?

they jump far compared to their body size–a beetle can't actually jump farther than a human (just think about a human long jump vs. a bug jumping–the bug just looks like it jumps far). Plus, different animals have different ways of getting around drag (here we're just making estimates for animals in general, not the few freak cases)

The long jump doesn't apply here - that is a running start very different from jumping from a standstill. My cat just lumped like 3 feet - and she has my jump height beat for sure.

they don't actually (if you and a beetle were in a jumping contest, you would win)

This fact comes intuitively to me and it's really interesting to think about this with calculations.

This is logical to me too.

I agree. It is definitely nice to be able to explain things that you see with mathematical models.

does this have to do with the drag coefficient at all?

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\mathrm{kg\,m^{-3}}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\mathrm{kg\,m^{-3}}$. The

Would it be useful to estimate it?

It's not a question of whether it would be useful (which it would be), but whether it is worth the trouble.

Here, Sanjoy cleverly evaded using a constant of proportionality by indirectly plugging in known numbers to get a ratio instead. Since it is a ratio, the constant of proportionality appears on both top and bottom, and cancels, eliminating the need to ever know what it was.

The ratio itself, tells us valuable information about "to what degree is the flea affected by drag", which is something that having a constant of proportionality would've also been able to do. The problem is that finding a constant of proportionality probably requires experimental data to "fit" the equation of interest, whereas doing the ratio here is a shortcut to getting some useful interpretation right away.

> Where you say "the constant of proportionality appears on top and bottom, and cancels" is false for any of the equations written on this page.
>
> What the equation above says, is that E=C/l, so E is proportional to 1/l. That constant is still there, hidden, and does not cancel. It is alluded to by the "proportional" sign.
>
> The constants do cancel, however, when you want to compare the magnitude of this ratio for a human to the magnitude of the ratio for a flea, for example. Since the human has a length scale of a thousand times that of a flea, this ratio, Edrag/Erequired, is a thousand times smaller for a human than for a flea. That doesn't give us any idea about its magnitude for a flea, unless we knew its magnitude for a human, and thus could derive C, the constant of proportionality.

why the difference (as in drag is no longer trivial in comparison to each) compared to the earlier drag lecture we had with the cones?

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\text{kg}\,\text{m}^{-3}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\text{kg}\,\text{m}^{-3}$. The

So basically this says that we forgot to put in the density of the animal into our original calculations, and this fixes it?

Not so much that we forgot, but that this is how one would go about using the proportionality calculations we ran through above into calculations of an actual limit (this is more of an equation, whereas the above derivation is getting us a relationship)

It almost seems like a way to throw out constants and other information to arrive at proportionalities and then somehow fix for the dropped things at the end...mildly hand-wavy but really cool

Where the top density is the density of our fluid, here air, right?

Oops, never mind. I see it.

I'm sure this works well for air. However, I feel that when I am in the water (say a swimming pool), I can jump higher despite the drag. Does that have to do with the boyancy of water?

Yes. It equates to a reduction in acceleration. While before we assumed that acceleration was a constant (g=9.8m/ss), in water your acceleration is a=F/m, but here F=|Wbody-Wwater|. Our analysis could have kept the variable for acceleration, and that would have explained your question.

Keep in mind water has a different density than air, of course, so the rho.fluid changes, too. I'm not sure if water's viscosity comes into play or not.

Yeah, it would also be interesting to see how the density of the environment affects the result- i.e. can you jump higher at sea level than on top of mt everest?

You'd be able to jump higher on top of Mt. Everest because lower rho in the numerator means lower drag energy ratio. (baseballs fly farther even in Denver, cyclists ride faster in Mexico City...)

Yeah! It gets even weirder than that. So many factors affect drag. The depth of a swimming pool, for example, can effect how fast you swim through the surface. That's so not-intuitive.

yeah, good point. let's just lump all that into 'edge effects' :)

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and ρAh is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator ρAh is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\text{kg}\,\text{m}^{-3}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\text{kg}\,\text{m}^{-3}$. The

I understand why this is important because of prior exposure to dimensional analysis. On the other hand, you may want to comment on why this is important in case the reader has not idea why dimensions matching matters.

This is a fair point - some readers may not understand why dimensions need to match.

I feel like if this is intended for an MIT audience it would not need to be explained. Also, I'm fairly sure the next unit is dimensional analysis, so this might be a sort of lead-in.

I agree, anyone who is reading through a book which analyzes the mass and drag of an animal should understand the concept of matching up dimensions.

I seem to remember that even in middle school and high school we would check answers based on if the dimensions matched... it certainly isn't a new concept for me that I"m learning in the class. It is cool to see how ti works in things like proportionality; it was cool how on the homework we got P Ev for example.

I think it is good that this line is here. It's like a good lead-in to the next unit.

I think dimensional analysis should have been included all along; it's really valuable as a "checking" method for your assumptions.

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\text{kg}\,\text{m}^{-3}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\text{kg}\,\text{m}^{-3}$. The

So the denser the animal, the less it is influenced by drag? Or do I have it backwards?

I am also confused by that.

yes. this is indicated by the animal's density appearing in the denominator. Think of tossing a baseball straight up into the air, and then a paper ball of the same size into the air with the same initial velocity. The baseball will go higher, and only because it is less influenced by drag.

i get it! that's so weird!

That's weird? Doesn't it make perfect sense?

I don't think there's a big range of animal densities, aside from cases where air volume is a significant fraction of internal volume. (Think of how you can change your buoyancy in water by having full or empty lungs.) I think for animals reducing cross-section is more effective.

To follow up, I found a classic book on this. Their values range from 0.99 to 1.1 specific gravity for mammals, and they mention the effect of air in the body. http://books.google.com/books?id=SV8rAAAAYAAJ&lpg=PA154&ots=wc5KfvUf

Whoa, I think this reasoning is flawed. If the baseball and paper have the same initial velocity, then the baseball will have more energy (making it go higher), since it has a higher mass. If the two balls have the same cross-section, mass and initial velocity, they'll behave the same way. Maybe if you tossed with equal energies this logic would hold up, since a less dense paper ball would lose more of its energy to drag, meaning it wouldn't go as high? The underlying physics of a tossed baseball and a jumping animal are different. We just went through a lot to find the relation of jump height to energy and length.

That's an interesting concept. I don't think it's fully accurate, as I would think drag is so much more influenced by sufrace area for instance that density itself might be negligible.

$$E_{drag} \sim \underbrace{\rho v^2 A}_{F_{drag}} \times h.$$

The ratio of these energies measures the importance of drag. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho v^2 A h}{m v^2} = \frac{\rho A h}{m}.$$

Since A is the cross-sectional area of the animal, Ah is the volume of air that it sweeps out in the jump, and $\rho A h$ is the mass of air swept out in the jump. So the relative importance of drag has a physical interpretation as a ratio of the mass of air displaced to the mass of the animal.

To find how this ratio depends on animal size, rewrite it in terms of the animal's side length l. In terms of side length, $A \sim l^2$ and $m \propto l^3$. What about the jump height h? The simplest analysis predicts that all animals have the same jump height, so $h \propto l^0$. Therefore the numerator $\rho A h$ is $\propto l^1$, the denominator m is $\propto l^3$, and

$$\frac{E_{drag}}{E_{required}} \propto \frac{l^2}{l^3} = l^{-1}.$$

So, small animals have a large ratio, meaning that drag affects the jumps of small animals more than it affects the jumps of large animals. The missing constant of proportionality means that we cannot say at what size an animal becomes 'small' for the purposes of drag. So the calculation so far cannot tell us whether fleas are included among the small animals.

The jump data, however, substitutes for the missing constant of proportionality. The ratio is

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho A h}{m} \sim \frac{\rho l^2 h}{\rho_{animal} l^3}.$$

It simplifies to

$$\frac{E_{drag}}{E_{required}} \sim \frac{\rho}{\rho_{animal}} \frac{h}{l}.$$

As a quick check, verify that the dimensions match. The left side is a ratio of energies, so it is dimensionless. The right side is the product of two dimensionless ratios, so it is also dimensionless. The dimensions match.

Now put in numbers. A density of air is $\rho \sim 1\,\text{kg}\,\text{m}^{-3}$. The density of an animal is roughly the density of water, so $\rho_{animal} \sim 10^3\,\text{kg}\,\text{m}^{-3}$. The

Sorry to correct your correction (9:58), but the baseball analogy is correct. Ignoring drag, a baseball and a paper ball will reach the same height because KE.initial=1/2mv^2 is equal to PE.final=mgh. The m's cancel and you get h=v^2/(2g).

But now add drag into the mix. Say the baseball loses 1/4 of its energy to drag, and the paper ball 1/2 of its energy. then the baseball will reach height 3h/4, but the paper ball will only reach a height of h/2.

Once they leave the ground, according to our model, animals and baseballs/paper balls behave in the same way.

> from 9:58 poster: You're right, I missed the m's cancelling, so that classic ball throwing problem from the first day of physics still works and vo determines h in the absence of drag.

Does this apply to humans as well?

Yes - I don't remember if the argument was made in this class or I saw it somewhere else, but humans are more than half made of water, so it's not a terrible rough estimate.

typical jump height – which is where the data substitutes for the constant of proportionality – is 60 cm or roughly 1 m. A flea's length is about 1 mm or $l \sim 10^{-3}$ m. So

$$\frac{E_{\text{drag}}}{E_{\text{required}}} \sim \frac{1\,\text{kg}\,\text{m}^{-3}}{10^{3}\,\text{kg}\,\text{m}^{-3}} \frac{1\,\text{m}}{10^{-3}\,\text{m}} \sim 1.$$

The ratio being unity means that if a flea would jump to 60 cm, overcoming drag would require roughly as much as energy as would the jump itself in vacuum.

Drag provides a plausible explanation for why fleas do not jump as high as the typical height to which larger animals jump.

Sometimes a summary table might be useful as a way to list the quantities that have been derived.

Completely agreed - I had to scroll back up the page to remember some.

This was also something I have been thinking about while reading some of the more numerical sections.

I agree - especially since it is hard for up to scroll between sections when reading on NB.

it's probably be cool to present, at the beginning or end of the book, a "common sense table" for obvious metrics (ex: time to fly to california, what 100km is like...etc). i know it makes more sense to internalize something that makes sense to us but it'd be cool just to have one anyway (kind of an extended version of what we had for our pretest)

it's probably be cool to present, at the beginning or end of the book, a "common sense table" for obvious metrics (ex: time to fly to california, what 100km is like...etc). i know it makes more sense to internalize something that makes sense to us but it'd be cool just to have one anyway (kind of an extended version of what we had for our pretest)

Are all these concepts useful to animal researchers in any way? I feel like this section is good in terms of analyzing drag once more, but in terms of practicality I find it a bit loose

I would like to see this same procedure applied to different animals so that we can compare the results.

typical jump height – which is where the data substitutes for the constant of proportionality – is 60 cm or roughly 1 m. A flea's length is about 1 mm or $l \sim 10^{-3}$ m. So

$$\frac{E_{drag}}{E_{required}} \sim \frac{1\,\mathrm{kg\,m^{-3}}}{10^3\,\mathrm{kg\,m^{-3}}}\,\frac{1\,\mathrm{m}}{10^{-3}\,\mathrm{m}} \sim 1.$$

The ratio being unity means that if a flea would jump to 60 cm, overcoming drag would require roughly as much as energy as would the jump itself in vacuum.

Drag provides a plausible explanation for why fleas do not jump as high as the typical height to which larger animals jump.

So does this mean that in order for a flea to jump 60cm, they would have to exert double the amount of force? (the force to accomplish the jump in a vacuum plus the force to overcome drag)

Good point. Does it scale linearly? I would guess yes, for our purposes, but I'm not sure.

I don't understand your question. If overcoming drag requires as much energy as the jump inside a vacuum, doesn't that imply that force needed is double the force within a vacuum.

Yeah I believe so. I had to reread it again to figure out that it didn't mean that the drag force was essentially zero. I think putting it in terms of doubling the amount of energy (like you mentioned) would have made it clearer.

yeah your comment helped me figure out what he actually meant by that.

It does definitely sound like he is saying that the drag force is zero. Perhaps it should be spelled out that a flea has to exert double the "vacuum force" in order to jump to 60 cm.

Yea I agree with the comment above. This would help clarify what was meant. I think just slight rewording would help like " if a flea would jump to 60cm, just overcoming the drag force requires the same amount of energy as performing the entire jump in a vacuum."

one too many "as" here

I think your explanation for this section makes sense.

Here, when you mention 60 cm again, I got a little confused. Perhaps make it more clear that 60 cm 1 m, and since we used 1 m in our equation, I think you should refer to it as 1 m the second time around.

Here, when you mention 60 cm again, I got a little confused. Perhaps make it more clear that 60 cm 1 m, and since we used 1 m in our equation, I think you should refer to it as 1 m the second time around.

At one point we assumed h to be independent of l, to continue in our calculation. Now that we know it's not true, it would be interesting to take a better estimate of h, and refine our estimate further.

typical jump height – which is where the data substitutes for the constant of proportionality – is 60 cm or roughly 1 m. A flea's length is about 1 mm or $l \sim 10^{-3}$ m. So

$$\frac{E_{drag}}{E_{required}} \sim \frac{1\,\text{kg}\,\text{m}^{-3}}{10^3\,\text{kg}\,\text{m}^{-3}}\,\frac{1\,\text{m}}{10^{-3}\,\text{m}} \sim 1.$$

The ratio being unity means that if a flea would jump to 60 cm, overcoming drag would require roughly as much as energy as would the jump itself in vacuum.

Drag provides a plausible explanation for why fleas do not jump as high as the typical height to which larger animals jump.

I will say that although we have looked at drag a few times in this course already, this one was the easiest example for me to understand.

I thought that flees could fly a little bit as well as just jump. How would you account for the animals that can jump, but can also fly, or sort of fly to help them jump further, such as a beetle, flying fish or flying squirrel. Is there any way to apply the same formulas to their movements?

I would think that since we're only considering height, gliding ability wouldn't come into effect.

Is there a way we can figure out how high they can jump by using this, or can we only say that they don't jump as far?

I feel like we've ignored a lot of constants along the route of proportionality so that the estimate would be quite rough.

true, we have been ignoring constants, but in finding proportionality, we don't need to worry about constants. for example, if we're finding the ratio of how high a human jumps vs a fly, we don't need constants that stay the same (like g) since they cancel out when we divide.

The above thread talks about how a flea would need double the energy that larger animals need since it needs energy to get to the height (mgh) and it needs an equal proportion of energy to overcome drag.

By this reasoning, fleas should be able to jump to half the height of larger animals, which seems a lot more realistic from what I remember in movies/other info.

I agree with this reasoning and the idea that its a rough guess based on dropping constants. It seems like some of the final proportionality we used (the ratio of densities) helped make this guess a bit more reasonable. I wonder if there is a way to combine divide and conquer and proportional reasoning to add in accounting for constants

wow Drag is indeed very important!

Thats amazing.

I definitely expected this result from the similarity to surface area to volume scaling.

I guess you were clever about it, I would definitely not have thought this was the case though. Great conclusion!

typical jump height – which is where the data substitutes for the constant of proportionality – is 60 cm or roughly 1 m. A flea's length is about 1 mm or $l \sim 10^{-3}$ m. So

$$\frac{E_{drag}}{E_{required}} \sim \frac{1\,kg\,m^{-3}}{10^3\,kg\,m^{-3}} \frac{1\,m}{10^{-3}\,m} \sim 1.$$

The ratio being unity means that if a flea would jump to 60 cm, overcoming drag would require roughly as much as energy as would the jump itself in vacuum.

Drag provides a plausible explanation for why fleas do not jump as high as the typical height to which larger animals jump.

Still, considering their size, it's impressive how high they can jump!

For sure, I think someone told me that if you look at the ratio of jump height/height of animal, a flea could jump three stories if it were the height of a human.

What about larger animals that don't jump, like bears or elephants?

They just weren't designed t jump. This analysis is only useful for similarly designed animals. This is similar to not comparing lift on a person to that of the lift on a plane.

I wonder if there is some optimal size for jumping? Clearly larger animals jump higher, as seen in this example, but at what point do they get too big and start to lose jump height? I think this would be interesting to explore.

Probably when they are large / heavy enough to hurt themselves when they jump - like he was talking about with an elephant - jumping would be enough pressure to crush its bones.

I wonder what other animals of similar size jump; would it be proportional, or are fleas just a special case?

So I completely spaced on doing this reading before Friday's lecture, but I find that doing it afterwards I get more out of it because you've already explained it in person once.

I like how this lecture goes from the most basic scenario, and adds in other factors one by one

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{\max} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for $v_{\max}$ to work out the ratio $v_{\text{swimmer}}/v_{\text{cyclist}}$, and then use that ratio along with v_{cyclist} to work out v_{swimmer}.

The speed $v_{\max}$ is

$$v_{\max} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{\text{swimmer}}}{v_{\text{cyclist}}} \sim \left(\frac{P_{\text{swimmer}}}{P_{\text{cyclist}}}\right)^{1/3} \times \left(\frac{\rho_{\text{swimmer}}}{\rho_{\text{cyclist}}}\right)^{-1/3} \times \left(\frac{A_{\text{swimmer}}}{A_{\text{cyclist}}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\text{m}\,\text{s}^{-1}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\text{m}\,\text{s}^{-1}$, significantly higher than the prediction of $1\,\text{m}\,\text{s}^{-1}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

What is A in this section? Not sure what this Area is in reference too, maybe i'll refer back to the previous reading.

It's the cross-sectional area. I should add a subscript to indicate that.

How did you make that conversion?

I am really enjoying this section. It's really cool to just use a ratio with something that is known to find something that is not.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

Comments on page 1

Read the 'swimming' and 'flying' subsections for the Sunday memo. (Don't mind this spurious box. It's just there so that I have a place to anchor the comment. Otherwise I have to make a global comment, which would pull up a whole separate panel when you visit the link.)

On Friday, you briefly went over notation used in this class (ie. tilde and alpha). In future drafts, you should mention that in the earlier sections.

After reading this, I feel like it belongs earlier in the chapter. It's a clear and easy example of proportional reasoning.

you mean in relation to drag?

yes, he is saying that you can compare how the drag on a cyclist limits its maximum velocity to how the drag on a swimmer will limit their maximum velocity. All you really have to change are the cross sectional areas and densities of the fluid.

I think it would be better to reference the actual section. For example, if you say "In section ___ we predicted the world record..." it would introduce the section better and let us know where to look if we need a refresher.

I agree...I just clicked around for 10 minutes trying to find that reading for a refresher

yeah, maybe perhaps briefly say what the variables are.

didn't we already do an example for the record for running? this doesn't seem too different.

shouldn't this be the same as cycling except there's a different drag?

what is P again?

power generated by the athlete

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

Wouldn't A change more in motion in swimming than cycling? Would this affect the accuracy of our analysis?

besides the obvious response that "it's close enough" i think you don't have to worry about it so much because the moving parts of your body in water are pushing through the water and generating less drag than your cross section.

Also, you are protected by the cube root – even an error of a factor of 2 turns into only 25% after taking the cube root.

Here being water

Yeah there should probably be a subscript: p_water.

I actually think it is intentionally ambiguous. That is, you use the density corresponding to the athlete's medium (p_water for swimming, p_air for cyclists).

p_fluid, then?

I think that indicating it's water makes the equation easier to conceptualize. From there, it is evident that a change in fluid requires a change in density.

what does "new" rho and A mean?

you would need to change the density to the density of water, and you would need to change the area – when a person is on a cycle, most of their body is against the wind, when you are swimming, it is a much smaller area! (since you are horizontal (or hopefully, more on your side) most of the time)

yeah "new" is perhaps just "different" here. it's clear that we're now talking about swimming.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

This seems quicker than doing the ratio method below...or is the point of the ratio method to be able to estimate everything instead of calculate?

I think the point is 1) to estimate and 2) its easier to do the ratios without a calculator than to do the calculation

I agree, and even if the numbers were nice enough that the multiplication was simple, taking the cube root without aid would be very painful.

I think another point of using the ratio method for this problem is to teach us a new method using a familiar example- that way, we can check our answer using a different technique. I think it's cool to see how 1 problem can be approached with so many different models!

Yes, I look at this as another example of all the different ways we've calculated things like miles per gallon. Multiple approaches can lead us to the correct answer.

This is rather similar to lecture Friday, where he pointed out that using the circle-area equation and always plugging in new numbers for the radius is actually the "long" way of doing things, even though it doesn't seem like it.

I find this to be one of the more enlightening things about this class. At MIT we have so many simple formulas plugged into our heads that it always seems easier to do the math. However, much of this relies on a calculator and complete information. In this case we have neither so we need to find a simpler method.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

It may be instructive, but it is really better? To me, I don't think you're saving any time doing it one compared to another.

I agree, I would prefer to do some complex calculations in my head than to go through all this math relating the 2 values. I certainly couldn't do all this in my head.

I agree as well, simple math is easier than comparisons when you know nothing, especially if the format is something we've done previously

Maybe that's why he said instructive instead of better.

Agreeing with all of these comments - it would probably be easier to just compute the answer directly. Perhaps we'll see why this method is instructive?

it's "instructive" because it's a relatively simple example to work though & you can verify the approximation by revisiting something we've done previously

I would say simple math is only easier than comparisons when the estimations are easier. For harder, more obscure estimations, I would rather compare it to something I am familiar with.

It is easier for the reason that you don't have to explicitly find the constant of proportionality. I'm not sure why he used and not proportional to, but there should be a constant of proportionality. That drops out when you take the ratio of v.swimmer/v.cyclist, and multiply by v.cyclist.

After having read the section, I have to disagree. Except maybe if I had a calculator handy, I would much rather estimate some simple ratios and multiply them than try to approximate a cube root in my head.

It's not always better, but it's aa additional skill that's very useful when the equations might be messier. It's good to have more than one way to think around a problem too, and this really is a different way of imagining the problem. What if you were trying to compare the price of two buildings, one 3 levels of 900 sq ft and one 5 levels of 800 sq ft. You wouldn't calculate a per square foot cost, you would just scale the cost of one by those ratios (3/5 and 900/800) to see if it was a better deal.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

Well, for starters, it never said it's better.

Besides, whether or not it's better really isn't the point nor is it relevant. Here is just ANOTHER method to do it, which is one of the first things we're taught in this course. There isn't necessarily a best way and this happens to be a different approach that will emphasize different concepts that we should understand.

Hence, it's instructive.

You can do these calculations even if you don't know the constants or even the exact equations, just how things (dimensionally or intuitively) relate to one another.

Or if you simply can't remember the constants off the top of your head.

i suppose the more options we see, we can choose which one makes the most sense to us, and it's beneficial in that sense.

To me (the instructor), it seemed obviously a better approach. Therefore, to answer this question I had to think quite a bit about the intuitive reasoning behind that bland statement of the "obvious".

The most fundamental reason, I think, is abstraction. The scaling argument in the text is implicitly building an abstraction: that top speeds are inversely proportional to fluid density^(-1/3) (and similarly for area). Theefore, do that computation once – use the abstraction – and then reuse it to get the similar result for swimming.

I have drawn a tree picture to show the abstraction, and I will put that into the textbook as well as show it in lecture today.

this reminds me of how, in 8.02, they'd always encourage us to keep our answers in symbolic form and then evaluate our answers based on if they made sense. (if A goes up, then B should go down. yes, it is at the bottom of the fraction...)

this requires prior knowledge of the cyclist problem, right?

Yeah, I think he assumes that we know the value of V_cyclist

Can we do the example of running as well, maybe really quickly in class?

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

Why repeat an identical equation? You should either remove it from above or simply reference it here.

I think it helps though given the next step in the math

I think it would be good to have the equation repeated if there is a subscript on the rho term. If the first was rho_water then it could be generalized here as rho_fluid

It's in a new section, it doesn't take much space, and there isn't any reason you should expect someone to flip back when you can just print it again.

It does seem out of place, although it is helpful and I can attest that I hate it when textbooks reference "the equation on page ___ and I have to go actually search for it again.

which area does this refer to?

cross-sectional area

In the drag equation, the area always refers to the cross-section that's directed towards the oncoming flow. In this case, the swimmer's A refers to their head, shoulders, and part of the arms and legs.

Maybe a diagram would show this better.

I forget, does the cyclist's cross sectional area refer to when legs are extended or not? maybe the different is minute due to the order of proportionality

we're looking at two different areas. one is the front of the person, one is the cross-sectional. i can see how this would be confusing.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

This makes it sound like its the swimmer and cyclist's densities, not the densities of their respective environmental fluid. I know it's not consistent, but I feel like p_water/p_air would be more clear and useful.

At its worst, someone could just read this formula without reading the context, and assume that it is the people's density (and reduce the ratio to 1).

I don't know, it seemed pretty obvious to me, and it did a better job of making a point of the general idea of what he is doing... Maybe density swimmer_fluid... or maybe just stating it explicitly?

I agree with the suggestion to change it to p_water and p_air, it makes it more clear and it's still obvious to which each is referring to.

It didn't seem as clear to me. However wouldn't the density of the swimmer be water and that of the cyclist be air? If so, wouldn't that be a better reference than the using swimmer and cyclist

I agree that it should be labeled as p_water/p_air because I thought it was referring to the densities of the swimmer and cyclist initially.

I agree. I'll revise it to use clearer subscripts.

shouldn't swimmer be on bottom and cyclist be on top for this part or am I missing something?

OP: oh, I just noticed the negative 1/3 in the exponent, my fault.

This combines divide and conquer with proportional analysis. I like the way this is broken down.

I agree–this method of breaking things down proportionally makes problems a lot easier. In homework problems, whenever there is multiplication we can break formulas into components like this

I like how this was set up in terms of proportions. Comparing swimmers and cyclists for each element is much easier than having to calculate the v-max for both of them separately.

Do you account for the power generated by seperate muscle groups? i.e. arms vs. legs? or do you just keep power constant?

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

I really liked the tree for this that you drew in class today. I think it would be really nice to have had that in the reading too.

I really liked the tree for this that you drew in class today. I think it would be really nice to have had that in the reading too.

Agreed - it helps clarify for me where different values are coming from.

I really liked the tree for this that you drew in class today. I think it would be really nice to have had that in the reading too.

Why don't we have to consider the actual bike ever? Whatever power the cyclist is generating is converted into his velocity through the bike, whereas the swimmer is only using his body, so shouldn't the cyclist be proportionally a lot faster anyway?

Not necessarily. Bicycles increase our output efficiency at the speeds at which they're used, but a swimmer does not really need the mechanical advantage of something like a bicycle to swim efficiently. It's basically because human power output varies with the speed of the motion – that's why we have bicycles to increase our speed and levers to increase our torque. Swimming, however, seems to fall closer to the peak output for humans (although some mechanical advantage, like fins, does seem to help).

does the bicycle itself play at all into energy or power?

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{\max} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for $v_{\max}$ to work out the ratio $v_{\text{swimmer}}/v_{\text{cyclist}}$, and then use that ratio along with v_{cyclist} to work out v_{swimmer}.

The speed $v_{\max}$ is

$$v_{\max} \sim \left(\frac{P_{\text{athlete}}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{\text{swimmer}}}{v_{\text{cyclist}}} \sim \left(\frac{P_{\text{swimmer}}}{P_{\text{cyclist}}}\right)^{1/3} \times \left(\frac{\rho_{\text{swimmer}}}{\rho_{\text{cyclist}}}\right)^{-1/3} \times \left(\frac{A_{\text{swimmer}}}{A_{\text{cyclist}}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\text{m}\,\text{s}^{-1}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\text{m}\,\text{s}^{-1}$, significantly higher than the prediction of $1\,\text{m}\,\text{s}^{-1}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

I know you have to assume some things to start the analysis, but I have a problem saying swimmers and cyclists generate the same power. It's plausible, but not that likely.

Cyclists produce energy principally through their legs, whereas swimmers use most of their body. I also fail to how they put out equal power.

I would be inclined to argue contrariwise. Legs are much more powerful than the arms and shoulders and cyclists are able to work with very efficient strokes. Swimmers are very limited in their range of motion with their legs. Anyhow, we are talking about differences in 20-30%, which are of little consequence here.

I agree with the previous comment - it's true that the powers generated are likely very different, but unless they're off by a factor of 10, we'll end up getting the correct order of magnitude answer in the end.

It sounds like someone who is a swimmer or cyclist got offended, to me...

When we consider power generated we should look at the environment in which each person is racing. The cyclist only has to deal with air resistance while the swimmer has to push against water. So it's true that the cyclist is using primarily his legs (which are stronger) and the swimmer is using his arms, the factors probably even out enough that we can assume both generate about the same amount of power.

Both athletes are using most of the muscles in their bodies. Have you seen a cyclist with a weak upper body? The fact is they are both working their body's to a maximum, which is why we can make this assumption about power.

i fail to see what you're getting at here. how the power is generated doesnt matter, only quantity. i agree that we're making some allowances here, and glossing over any discrepancies.

Being that these are two specimens of simnilar physical fitness of the same species, it's reasonable to say they generate the same power within an APPROXIMATE range.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

Isn't the answer 1/10 of the speed of a cyclist that was calculated for as the average speed over a whole hour? If you are trying to calculate the maximum almost-instantaneous swimming speed you should not compare it to a cycling speed of such different time difference. And if I happen to be confused then you should be more explicit about what speed ratio you are using, more specifically what speed for the bicycle you are using.

I commented on this in the cycling reading, but I was bothered then when we didn't worry about time. We seemed to be calculating a maximum, but then we arbitrarily assumed this power could be applied for an hour. I'm not really sure what we're looking for with the swimming speed, instantaneous or over some period of time.

why would you say this when clearly there is a third factor? of course any formula is going to be incorrect if you're leaving things out.

I don't think it's to point out that we need another factor, instead it helps us understand how the final factor is going to affect our answer. Without this third factor, we are well off.

Here I had to go look up our results from a few readings ago to see how you arrived at this. Would it be possible to repeat somewhere in here that we had originally estimated v_cyclist = 10m/s?

Good point, I'd like to see that too.

I agree. also, what swimming world record? because for the cyclist, the speed we found was the speed over 1 hour...

I think review like this is good for an online textbook where we read different sites/chapters days apart , but not necessary for a published book where it's possible to flip through pages

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\text{m}\,\text{s}^{-1}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\text{m}\,\text{s}^{-1}$, significantly higher than the prediction of $1\,\text{m}\,\text{s}^{-1}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

They swim much faster in a 50m race than a mile race! I think the 50m record is like 20.91s, making the speed 2.39m/s.

Sprinting is much more entertaining but I think it might be the distance in a bike race and the mile swim might make the type of muscles used more comparable. (powerful fast twitch vs. endurance)

Right, we need to keep in mind our comparison. Can't compare distance/endurance with sprinting.

Besides, it makes little sense to compare burst of speed competitions to a more settled rhythm.

It's be interesting to compare the records for cycling in a sprint to a sprint swimming. Though it's a relatively long sprint (compared to 50m which I can't find anything on cycling for), the record for 200m cycling is 11 seconds. That's 18 m/s! Which is still significantly higher than the 2.39 m/s mentioned earlier for swimming.

It's ok that we used the distance swimming race here for our comparison because the speed we got was still much faster than what we expected.

It took me a few seconds longer than it should have to read this. Could you maybe put a space in so it's more clear? I was expecting a speed since that is what we are actually interested in.

also, maybe write out 'min' or 'minute', since m is used as meters everywhere else. Or just convert the number to seconds for us.

If we are talking about the world-record fastest swimmer, then why would you take a 1,500m race, instead of a 100m race?

The argument was made in another comment thread that the cyclist power was based on endurance, so he uses an endurance swim race as well.

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

I believe this is the only long-course Olympic record that predates 2008 (and Wikipedia backs this up).

Interesting...thought this would have been broken more recently with all the fuss about how improvements in swimming technology have made athletes faster.

This season the NCAA banded most of the really high tech suits but you would think that the high tech suits would have really helped especially in a distance event

Improvements in swimming technology? What makes a high tech suit...?

I was curious, so I looked into this as well. Apparently, high tech suits usually cover more of the body, and because they are "corset-like," enable those with stocky and muscled bodies to be as streamlined as a long and lean one, and a soft abdomen as effective as six-pack abs!

Where does this discrepancy come from?

why are we looking at how close a third of our process is? doesn't it make more sense to wait and compare the final solution anyway?

it may gives a us a better idea of where our estimations should lie if we didn't have any clue about the area. but then we would need the answer in the first place. yea this seems strange.

how should we know how close our prediction has to be? usually within a power of ten is supposed to be good enough, which this is

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

When I think about a swimmer traveling 1 meter in one second versus almost 2 meters in one second, I can see how 1.7 m/s is a lot higher than 1 m/s. However, given that this class is always working on the scale of powers of 10, my first reaction was that these two velocities were very similar (almost equal to 1). Why is it that such a small difference matters here? Is it because the numbers are so small?

I agree with the above, its large given what we know about swimmers, but given that we just nonchalently drop factors of two everywhere, I don't see how we can come to this conclusively.

This is only off by about 70%. In class we definitely would have said that this is a very good estimate. We drop coefficients of 2 all of the time. So why is this 1.7 coefficient suddenly significant?

I believe this is "significantly higher" due to the low order of magnitude.

The difference between 30m/s and 30.713m/s is, relatively, a lot smaller than the difference between 1m/s and 1.713m/s.

In general, there is no silver bullet for what is close enough and what isn't. We have to make that decision based on the context.

I agree with most of the statements above... while it it is a much smaller order of magnitude, so the percent difference is much larger, I'm a bit skeptical given how many other assumptions we're making that we wouldn't have just gotten some other closer speed by accident.

maybe stating it as miles per hour would be helpful. Then it is more apparent that the speed is almost doubled. 2.2 mph to 3.8 mph.

Maybe briefly mentioning mph just to emphasize the difference would be nice but then again it is another conversion that will lost more accuracy and is much less convenient than the conventional units we used here

4.4.2 Swimming

A previous section's analysis of cycling helps predict the world-record speed for swimming.

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

To evaluate the maximum speed for swimming, one could put in a new ρ and A directly into that formula. However, that method replicates the work of multiplying, dividing, and cube-rooting the various values.

Instead it is instructive to scale the numerical result for cycling by looking at how the maximum speed depends on the parameters of the situation. In other words, I'll use the formula for v_{max} to work out the ratio $v_{swimmer}/v_{cyclist}$, and then use that ratio along with $v_{cyclist}$ to work out $v_{swimmer}$.

The speed v_{max} is

$$v_{max} \sim \left(\frac{P_{athlete}}{\rho A}\right)^{1/3}.$$

So the ratio of swimming and cycling speeds is

$$\frac{v_{swimmer}}{v_{cyclist}} \sim \left(\frac{P_{swimmer}}{P_{cyclist}}\right)^{1/3} \times \left(\frac{\rho_{swimmer}}{\rho_{cyclist}}\right)^{-1/3} \times \left(\frac{A_{swimmer}}{A_{cyclist}}\right)^{-1/3}.$$

Estimate each factor in turn. The first factor accounts for the relative athletic prowess of swimmers and cyclists. Let's assume that they generate equal amounts of power; then the first factor is unity. The second factor accounts for the differing density of the mediums in which each athlete moves. Roughly, water is 1000 times denser than air. So the second factor contributes a factor of 0.1 to the speed ratio. If the only factors were the first two, then the swimming world record would be about $1\,\mathrm{m\,s^{-1}}$.

Let's compare with reality. The actual world record for a 1500-m freestyle (in a 50-m pool) is 14m34.56s set in July 2001 by Grant Hackett. That speed is $1.713\,\mathrm{m\,s^{-1}}$, significantly higher than the prediction of $1\,\mathrm{m\,s^{-1}}$.

The third factor comes to the rescue by accounting for the relative profile of a cyclist and a swimmer. A swimmer and a cyclist probably have the same width, but the swimmer's height (depth in the water) is perhaps

I think a factor here is that while a person is cycling, the wind / drag will ALWAYS be against the athlete. No amount of power will be able to significantly change the volume of air around the athlete, nor it's direction (only perhaps the small volume behind the cyclist, but that doesn't help him/her).

In contrast, swimmers are in a pool with a limited amount of water. Compared to the volume of a swimmer, it's not much. If you consider all the swimmers in each lane, the water:swimmer ratio is even less. In swimming, the athletes actually cause the water to move one direction while they go, so for the length of a pool, they're actually aided by the tides/waves they create. This could be in part responsible for this disparity.

Cyclists can also be aided by a tailwind. If the athlete happens to be cycling in a windy area with a very strong tail wind, the wind could also aid them.

a profile figure might be nice for this

This isn't actually the swimmer's height is it? I thought we were measuring the cross sectional area going against the water?

That's what he's saying. The swimmers height is defined by his depth in the water, I'm sure this assumes that the swimmer isn't completely submerged which is usually true in a swimming race.

I think this makes sense–I mean either way, whether or not he is fully submerged or not, the cross sectional area is significantly less: in the swimming case, you're looking at basically a head's height and shoulders' width, whereas for the cyclist you're looking at shoulder's width over a much longer height...

haha unless you've seen a butterfly swimmers shoulders.

The order in which you made this analysis was a bit off-putting. It felt like you might have somehow made the numbers work, instead of allowing them to naturally fall out correctly.

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\mathrm{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\mathrm{plane}}}{E_{\mathrm{car}}} \sim \frac{\rho_{\mathrm{up-high}}}{\rho_{\mathrm{low}}} \times \left(\frac{v_{\mathrm{plane}}}{v_{\mathrm{car}}}\right)^2 \times \frac{A_{\mathrm{plane}}}{A_{\mathrm{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\mathrm{plane}} \sim 600$ mph and $v_{\mathrm{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

it just seems to me that making the guess of one-sixth after the fact that the first estimate was too high seems a bit contrived?

I agree - it seems like you just picked a number that would make it close.

yeah i agree...is there a way to make this approximation more accurately so that it's not just a somewhat random number?

I think it feels contrived because it is so close. But that is the power of good guesses. After all, the world obeys the laws of physics so why shouldn't we get a close answer, if we follow the physics relations as closely as we deem possible. I just laid on the floor to measure my height and it does come out to be about 1/6th of how high my bike is.

I find how close you get to real values to be quite remarkable given how many assumptions you make.

I think that it was deceptive of you to compare your estimate with the speed of swimmers in a 1,500m race, since they are pacing themselves. This estimate is most likely not very accurate.

I think that it's well within the error of the reasoning we've used to get this far. I don't think it's too much of a problem.

i agree. we have to consider the timescale over which these things are happening. the 50m sprint is likely what we're looking for here.

I think it was unclear but not unfair. First, I reorganized the sections badly, and somehow dropped the analysis of cycling (so the "previous section" points to nowhere).

Second, that analysis, which you had no way to know, compared the prediction to the record for long-distance cycling. In particular, the record in question is the "one-hour" record: How far can one cycle in one hour? That is a fair comparison with a long-distance swimming record.

For cycling, the record one-hour distance is about 50 km, for a speed of 13 or 14 m/s, which I took as reasonable agreement with the prediction of 10 m/s.

never would have thought of that

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

Cool! This is very close to Michael Phelps record for the 200m freestyle- 1:43.86 which comes out to be 1.9 m/s

Dang he's fast...

does this mean swimmers have come close to "maxing" out the record?

They can swim much faster than 1.8m/s, the number from above is for a mile – and you can imagine that one would swim much slower over the coarse of a mile – the shortest race these days is 50m (for a long coarse pool), and I think I mentioned above that is about 2.39m/s.

but they have come close to the max _for long distance_ swimming.

There isn't a hard upper bound. Reducing their drag coefficient, changing swimming style, increasing potential power output, etc., could all increase maximum speed. This is just an estimate.

this is why swimmers are relying on technology, mainly new body suits to make them swim faster

Definitely true, the new suits they use significantly reduces drag for the swimmers in the water. This is still an impressively close calculation!

This isn't a maximum or even an estimate of one. The number relies directly on the athlete's power output, so better training may allow them to increase that number and exceed our estimate.

To me it kind of seems way too convenient that the numbers worked out the way that they did.

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

But we still needed to know the maximum velocity of the cyclist which we approximated and worked out using arithmetic. So I guess I would note that one downfall is that you'll have to use raw approximations without scaling for something, and then use that as the base for subsequent approximations.

I think this is just another example of proportional reasoning and how scaling can save us time

I'm not sure that the scaling was too much faster than plugging in new numbers would've been - we remember P, know the new rho and can guess the new A...things we have to do anyways to create the scaling factor. Is there some error propagation benefit to scaling?

One reason that it's beneficial to scale is to easily compare your estimate for one number with that of a number you already know. It makes it easier to see if any of your estimates are unrealistic.

I think this example doesn't really seem that much better with scaling because it doesn't allow us to eliminate anything - we just had to do the whole thing out basically. Maybe taking an example where there were constants in the formula or larger similarities - say between the areas - that would allow us to eliminate numbers by using the scaling trick.

I just realized, in writing another response, that I forgot to include the cycling calculation when I reorganized the book draft. Sorry about that.

The main results from the ghost section are the prediction of 10 m/s and the world record of about 13.5 m/s (for the maximum speed over a one-hour stretch).

So the purpose of kind of re-doing this example was to understand the source of estimation's accuracy?

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

This is a very good point. We could hung up on details such as whether or not the cyclist and the swimmer actually exert the same amount of power. But the debatable issue decreases in significance after cube rooting.

The opposite is also true though, and I feel like he glosses over the fact. If something goes like a factor squared or cubed, minor errors become greatly exaggerated.

Yeah, this is true, actually. while in this case the cube roots make the errors smaller in the final answer, there are cases where things are squared, and error propagates much larger.

This part does clear up my question.

Correct me if I'm wrong, but we never made an estimate for power produced, simply assumed they were equivalent values (as it should be, if it's the same muscles being used). Perhaps you meant about the densities of water?

I think this might be referring to the equation for velocity that we didn't bother using; even if we didn't estimate the power directly, the cube root compresses the P term. It probably didn't mean the densities because the values for densities of water/air are more or less correct, and the only things we're estimating are A and P.

This does make it seem like we estimated the power (but even then this doesn't specify whose power this is referring to). We should either do the estimate earlier or nix this specific value.

I think we made the estimates of power earlier in the cycling example–remember the stair climbing/time calculation? You're right, we don't have to think about them again here since they're not changing, but if we were to compare people capable of different power outputs (due to age, fitness, etc.) we would need to do more estimation for the ratio above.

It does say "for example" right before, giving the impression that it is all hypothetical. I guess it could be fixed by writing "an example" instead of the "the example."

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

ah yeah, I need to start thinking about this more – its a bit confusing! So you do the factor you might be off ("n") to whatever power it is, so n^1/3?

Try it with some sample numbers–say you had an estimate of the volume of a cube and wanted to find its radius, any errors you made in the volume would be scaled down by the cube root. If the true volume were 1000 (length = 10) and you guessed 1100 (10% off, a factor of 1.1) you would get 10.3 (only 3% off, since 1.03 = the cube root of the error factor of 1.1). Of course, there are cases that magnify the error, for instance, going from radius to volume, the reverse of my example. Lesson: be more careful with your original estimate if your error is scaled up, or at least be aware that it is scaled up.

This may be a bit philosophical, but what is a "mistake" in estimation? In other words, at what point do we say a number/answer is "wrong?"

I think "mistake" is a bit harsh here, maybe misjudgement

Yeah agreed, "mistake" is not the right word to use here.

Whether or not it was a mistake to use the word mistake or not, I think that the point here is pretty clear.

I think if something isn't exact it's always 'wrong'. There are just varying degrees of wrongness. For example. This was wrong by 25%. I think a good way of judging if something is "too wrong" to be a good estimate is if it is outside of an order of magnitude estimation.

I think if something isn't exact it's always 'wrong'. There are just varying degrees of wrongness. For example. This was wrong by 25%. I think a good way of judging if something is "too wrong" to be a good estimate is if it is outside of an order of magnitude estimation.

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

this seems to happen in most cases lol (at least so far)

If you estimate enough numbers, it is probable that some will be overestimates and some underestimates. I'm not sure whether its really "luck"

That's a good point, I never looked at the canceling that way.

It would be nice to see a more rigorous explanation of how things always seem to cancel, because right now it seems they just do. It doesn't have to be a mathematical proof or anything, just a general "well, we overestimate on things in both the numerator and denominator and it works out well" would be fine.

I agree, I don't remember talking about error analysis much in the readings... this is cool!

It is cool. However, do you think he plans to over estimate and underetimate in every case, or is it luck?

I don't usually get that lucky in my estimation mistakes.

me either. are you picking problems where errors cancel each other out on purpose? if you had two errors that contributed in one direction by 25% each...that would make your answer pretty inaccurate. I think a 25% error rate is already pretty bad...but I guess we are usually only looking at orders of magnitude

I think he just consciously makes a point to remember when a value has been overestimated, and thus will underestimate another to compensate.

I definitely do that. If you know one of your estimates is an overestimate, why not consciously look for an underestimate elsewhere to balance the two?

Is there some way for us to intentionally increase the chances that our errors will cancel?

Yeah, what if our errors compounded into more egregious ones?

I think that's why we always have to check our answers. If we're dealing with something where you can't check your answer, assume the worst with your errors and have wide confidence intervals.

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\mathrm{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\mathrm{plane}}}{E_{\mathrm{car}}} \sim \frac{\rho_{\mathrm{up-high}}}{\rho_{\mathrm{low}}} \times \left(\frac{v_{\mathrm{plane}}}{v_{\mathrm{car}}}\right)^2 \times \frac{A_{\mathrm{plane}}}{A_{\mathrm{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\mathrm{plane}} \sim 600\,\mathrm{mph}$ and $v_{\mathrm{car}} \sim 60\,\mathrm{mph}$, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

not really understanding this error analysis..

it's just saying to utilize the powers of the ratios here in order to reduce order of magnitudes. i'm not sure its something you could account for at the beginning though.

I think it's more about understanding how close your approximations really have to be. We are given a bit of leeway in this problem because we are taking the cube root so it is not so important to be perfectly accurate.

It would help to say this - I saw in another comment a reply about how a factor of 2 becomes only 25%. Including that here would no doubt clarify matters.

I feel pretty familiar with drag now.

what does it mean to "scale" the drag formula? i thought it was just ready to receive values of various magnitudes, depending on what we were working on.... what scaling is needed?

agreed, I think the drag coefficient takes care of this.

This is probably dumb, but I don't really know the difference between a jumbo jet and a 747 off the top of my head. Obviously I can look it up but its not something I know.

As far as I know they're different names for the same thing.

I don't think the difference is significant. Basically the message is: a really freaking big plane!

Speaking in terms of sets, a 747 is a subset of jumbo planes, probably the most well-known one :)

I wonder why this would be seen as large and meaningless. couldnt we work backwards in the future and using the fuel efficiency of jumbo jet to find that of a plane and car

I'm sure we could, but it's much easier for people to understand if we find the fuel efficiency of a jet in terms of that of a car–something we are all familiar with.

Agreed. The point about it being large is that we should look at things that we come across in our every day lives.

is this different that a 747?

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\mathrm{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\mathrm{plane}}}{E_{\mathrm{car}}} \sim \frac{\rho_{\mathrm{up-high}}}{\rho_{\mathrm{low}}} \times \left(\frac{v_{\mathrm{plane}}}{v_{\mathrm{car}}}\right)^2 \times \frac{A_{\mathrm{plane}}}{A_{\mathrm{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\mathrm{plane}} \sim 600$ mph and $v_{\mathrm{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

I love that this plane fuel consumption problem keeps coming up, and that we've had the opportunity to solve it different ways each time. More then any single example, I feel like this one has been a good exercise for each method.

Yes definitely agree! Having one common thread like this through different estimation methods really highlights strengths and weaknesses between the diferrent types.

I have noticed that this is a very common way of finding solutions, but wouldn't you want to use an independent way just as a check for your answers?

why is it meaningless? we've been coming up with a lot of large numbers.

I like this explanation - it makes scaling seem much more reasonable because it produces an answer that is more easy to think about for the average person.

This is a profound statement. It's interesting to consider the amount of fuel required to bring rockets to orbit vs. having them coast in space where no atmosphere exists and no extra boosts are needed

i never knew that drag was so important in all these applications, the amount of places you can do a drag analysis is amazing

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

This is helpful so I know that the actual complexity of the problem will be explain later, so I don't have to ask now.

Agreed, this is definitely an improvement from some of the previous articles.

Sorry to disagree here, but I think stating this takes away from the strength of your methods and approximations. Discuss limitations and further refinements after you have produced an initial prediction. Simply stating that you are are making this assumption should suffice for any doubters until you get to this discussion at the end of the section.

Disagree – if you just look at comments from earlier sections, a lot of people got caught up early on when he made some assumption people found hard to believe (i.e., cubic mountains). So it helps to know that it'll be explained – because at the very first half of the sentence I was already skeptical that most of the jet fuel goes to fight drag.

I agree that it's useful to point out there are further complexities not yet covered, but I also agree that it shouldn't sound like the estimate is doomed from the beginning. Perhaps saying that "the assumption is not quite right - but still allows us to arrive at a reasonable estimate" would be helpful?

I agree with the fourth post as well. A comment like that really brings more perspective into the problem as you are doing it. In a way it makes me feel somewhat more confident, even though the "assumption is not quite right"

I totally agree - as soon as i read this I thought - oh that's nice it gets rid of all the comments regarding shortcomings that are eventually explained anyway!

as opposed to having a drag coefficient that goes towards geometry?

I think I've finally got this one memorized!

For some reason I keep forgetting all of these.

i think they'll become easier once we hit the dimensional analysis chapter! :)

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\mathrm{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\mathrm{plane}}}{E_{\mathrm{car}}} \sim \frac{\rho_{\mathrm{up-high}}}{\rho_{\mathrm{low}}} \times \left(\frac{v_{\mathrm{plane}}}{v_{\mathrm{car}}}\right)^2 \times \frac{A_{\mathrm{plane}}}{A_{\mathrm{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\mathrm{plane}} \sim 600$ mph and $v_{\mathrm{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

when did we switch from fuel efficiency to fuel consumption?

probably because

fuel efficiency = energy consumed in a distance d

fuel efficiency = E/d

since we compare the same distance for both plane and car, the ratio of their respective fuel efficiencies is simply the ratio of their energy consumed.

a little confused on what this term means..

Why "up-high", but not "down-low"?

Because then you'd have to fit in a too-slow

This is to refer to the density of the air at a higher altitude (I think). The 'down low' is rho low underneath, which is roughly the density of air under normal conditions.

Shouldn't you be accounting for the amount of people each can hold? A car that holds 5 people can not be compared to a plane that holds 300.

Sorry I noticed your discussion on the passengers after the first note was written.

you could always delete your note, if you want... I think nb still saves a "hidden" deleted note which the admin can see.

(one reason for deleting a note you don't want anymore is that it cleans up the clutter, since everyone draws boxes in the same place)

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

much better than 's'

So much for using 'r' for 'range' instead...

Not that this matters, but perhaps it should be d_plane/d_car.

i thought 'x' would have been better than 'r' anyway...

x sometimes becomes a problem when you are also using it as the multiplication sign

Clearly the appropriate choice of variable to represent distance is a contentious issue... Deep breaths guys, it's really not a big deal.

You're right, you get the picture whether its d, s, r, x or whatever... as long as it's defined. I agree with above though, that it might be better to use the plane and car subscript with the distance variable.

And d_plane and d_car are the same thing. There is no reason to put any more detail than what's already up there.

I'd still like to have seen the intervening equation–E=Fdrag*d since we both added that and took the car/plane ratio in one step.

Based on the analysis in the preceding section, can we assume that the squared factor would amplify any errors in estimations of velocity?

Yes, it would. but in this case, this value is probably the easiest to estimate–the relative speeds looks much easier to guage (for me) than relative cross-sectional areas or air densities at various elevations...so even though error would amplify, it probably shouldnt be too big a deal.

But what about holding a plane up in the air?

Never mind, it's mentioned later on.

I like how you break everything down into manageable estimates. This makes the problem much less difficult, and it really proves your points in earlier memo's about using abstraction to your advantage.

I agree...these examples that make use of the previous techniques we've learned in conjunction with the new one we're studying really help bring things together

one-sixth that of a crouched cyclist. So the third factor contributes $6^{1/3}$ to the predicted speed, making it $1.8\,\mathrm{m\,s^{-1}}$.

This prediction is close to the actual record, closer to reality than one might expect given the approximations in the physics, the values, and the arithmetic. However, the accuracy is a result of the form of the estimate, that the maximum speed is proportional to the cube root of the athlete's power and the inverse cube root of the cross-sectional area. Errors in either the power or area get compressed by the cube root. For example, the estimate of 500 W might easily be in error by a factor of 2 in either direction. The resulting error in the maximum speed is $2^{1/3}$ or 1.25, an error of only 25%. The cross-sectional area of a swimmer might be in error by a factor of 2 as well, and this mistake would contribute only a 25% error to the maximum speed. [With luck, the two errors would cancel!]

4.4.3 Flying

In the next example, I scale the drag formula to estimate the fuel efficiency of a jumbo jet. Rather than estimating the actual fuel consumption, which would produce a large, meaningless number, it is more instructive to estimate the relative fuel efficiency of a plane and a car.

Assume that jet fuel goes mostly to fighting drag. This assumption is not quite right, so at the end I'll discuss it and other troubles in the analysis. The next step is to assume that the drag force for a plane is given by the same formula as for a car:

$$F_{\text{drag}} \sim \rho v^2 A.$$

Then the ratio of energy consumed in travelling a distance d is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{\rho_{\text{up-high}}}{\rho_{\text{low}}} \times \left(\frac{v_{\text{plane}}}{v_{\text{car}}}\right)^2 \times \frac{A_{\text{plane}}}{A_{\text{car}}} \times \frac{d}{d}.$$

Estimate each factor in turn. The first factor accounts for the lower air density at a plane's cruising altitude. At 10 km, the density is roughly one-third of the sea-level density, so the first factor contributes 1/3. The second factor accounts for the faster speed of a plane. Perhaps $v_{\text{plane}} \sim 600$ mph and $v_{\text{car}} \sim 60$ mph, so the second factor contributes a factor of 100. The third factor accounts for the greater cross-sectional area of the plane. As a reasonable estimate

This would have been helpful information for the last pset (the last problem on the pset required this information).

I think this was brought up in lecture before pset #4 was due, but not sure when exactly.

The monday before, I think. It was still something you could estimate, though, if only with a gut feeling.

I think it would be useful to put the e^-(h/H) formula that you showed in class into the text. I thought that was pretty handy.

You noted in the previous section about how an error of a factor of 2 in power estimation only came out to a factor of 1.25 in our result. It might be good to show the dark side, too, and note that any error here gets squared (or doubled, depending on how you think about it... 10% error goes to 21%...).

I think our 600 and 60 numbers are relatively accurate, so this may not apply here, but at least noting it somewhere would be useful and honest.

Yeah I agree, as I read it I was already asking myself what would happen if the errors didn't cancel out but only compounded on one another, so a "dark side" problem would be useful.

$A_{plane} \sim 6\,m \times 6\,m = 36\,m^2,$

whereas

$A_{car} \sim 2\,m \times 1.5\,m = 3\,m^2,$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{plane}}{E_{car}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{drag} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

Comments on page 3

Is this an average area? the nose is more narrow than the middle part of the plane, where the carry on baggage also contributes area.

When you are computing drag you only care about the frontal cross sectional area (The entire area you can see from the front) so a tapered nose doesn't affect the area, only the coefficient of drag which we left out of the Fdrag equation back in reading 12 or so

why do we even need this? (it's just equal to 1, I don't see why we need it)

the fourth factor is the ratios of the distances, right? since energy = force * distance, you need the distance term...

I think it's important so we compare on the more familiar basis of energy, but it was an unstated leap from F to E above.

I was wondering this as well.

I think its just for consistency, so we keep track of what we are calculating.

I think it's good to see any unities done out. I'd rather have them there so it's understandable where it came from.

Is it unity though? A plane can go straight across the US while a car has to follow roads which add about 400 miles to the trip (2600 mi flying vs 3000 driving). It's a difference of about 15%, is that small enough that we can just discount it?

By saying that it contributes unity, this example is assuming that the plane and car travel the same distance.

How much would this actually effect the overall result?

Whenever these planes hypothetically fly to Los Angeles it makes me homesick.

hahaha, me too

when comparing energy efficiency, it is important to include the inefficiency of indirect freeways, and the requirement to stop to refuel, sleep, go to the bathroom. Also traffic, and elevation gain and loss play a role. I would guess flying cross country is far more efficient than driving.

$$A_{plane} \sim 6\,m \times 6\,m = 36\,m^2,$$

whereas

$$A_{car} \sim 2\,m \times 1.5\,m = 3\,m^2,$$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{plane}}{E_{car}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{drag} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

I really liked this example in class.

Yeah! I also like how he waits until the end to account for passengers, it makes you think after the first calculation why it's so off from what you'd expect.

I think the entire idea of looking at a problem and how each factor contributes to the answer is amazing. Even though we could do this stuff with the simple equation, it is much easier to relate it to something we understand, like cars.

...but not a bus

But there are a lot of other perks that make flying much more appealing to many people like the time factor. Are the factor of time and perhaps the specialization of pilots the main factors in price difference of travel by the two modes (since you've said fuel efficiency is pretty similar as well)?

Another thing to think about is that planes are going to make their trip to California with or without you, full or not. By choosing to drive instead of taking a plane you are in a way being more wasteful by inaction as well as what you use driving.

Haha, very true!

Even in MA. I noticed this the other day when I was stuck in traffic during rush hour.

I thought this was funny, but wondered why he chose California when it hasn't been mentioned anywhere else? I am from there, so I can at least admit that it IS quite fuel inefficient by and large...

Sanjoy got degrees from Caltech and Stanford–that would be a non-negligible # of years in California.

http://web.mit.edu/tll/about-tll/mahajan.html

typo

does one or the other provide a more "green" environmentally friendly option? or are they about the same?

$$A_{plane} \sim 6\,m \times 6\,m = 36\,m^2,$$

whereas

$$A_{car} \sim 2\,m \times 1.5\,m = 3\,m^2,$$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{plane}}{E_{car}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{drag} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

Found some cool data. As of 2001, this number was 1.57 for cars across the US (and higher for SUVs/Vans, lower for trucks and motorcycles)

http://www1.eere.energy.gov/vehiclesandfuels/facts/2003/fcvt_fotw257.html

Again, I still think the number would be even higher for the average trip to Califonia.

But are they in fact really equal? I read below that the 2 errors in estimating scaling ratios for the plane basically compensate for each other, but in general are the relative efficiencies about equal? I would have guessed differently.

This seems reasonable, but I would like to see the facts here to compare. (Or I could just look it up again.)

Didn't we figure this out in class on wednesday?

I believe we did. I'm wondering why we did it in lecture before seeing it in the reading. It seems backwards from how we've done everything else. Perhaps he forgot he also used this example in the reading?

$$A_{plane} \sim 6\,m \times 6\,m = 36\,m^2,$$

whereas

$$A_{car} \sim 2\,m \times 1.5\,m = 3\,m^2,$$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{plane}}{E_{car}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{drag} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

I would think this number would be higher, especially with the focus on environmental responsibility and carpooling.

I agree, I think that a more accurate approximation would be that a typical car carries 1-2 people, so I'd use 1.5 people. This subsequently causes the fuel efficiency of a plane and the fuel efficiency of car to be unequal, which I think seems more plausible.

in lecture he said something about actually sitting and watching cars ... it seemed to be a good thing to go with from that.

from my experience watching cars, i would actually guess 1.1 or 1.2 ... yes, the typically carry 1-2 people, however they carry 1 _a lot_ more than they carry 2 (or more).

Yeah but how many people drive to california alone? It's usually a road trip of some sorts. I've never heard of anyone driving out west with fewer than 2 people in the car.

Well, the distance is just an arbitrary distance that scales to unity in the ratio anyhow, so the fact it's a road trip is irrelevant. Generally for every day use of a car, people are driving to and from work and it does tend to be about one person to a car...

The average certainly can't be one person per car because there aren't any cars with less than one person, so I agree the estimate should be higher. Is there a reason you talk about the typical car instead of the average number of people per car?

well, with the point being that a jumbo jet capacity is 400, does 1 or 2 or 3 people really make that much of a difference? and I agree with the 6th reply that the fact that its a road trip is irrelevant also.

Agreed. Just remember, it's an approximation class!

That is a really interesting result. However, I would think that a plane would be more efficient given the fact that they are strongly preferred for long distance travel.

This answer really surprised me too–but I guess it still makes sense that planes would be preferred–plane trips are shorter, and it's less of a hassle when you can sleep and don't have to drive. But also, this estimate is fuel efficiency based on drag alone. There are many other ways planes and cars waste energy right? do things like AC play a significant role?

$$A_{plane} \sim 6\,m \times 6\,m = 36\,m^2,$$

whereas

$$A_{car} \sim 2\,m \times 1.5\,m = 3\,m^2,$$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{plane}}{E_{car}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{drag} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

is it the same kind of fuel that is used by both the car and the plane? I'm assuming one is more cheaper than the other...

that's wicked cool

When you say it is comparable, do you mean that it is about the same amount of energy? Because wouldn't that mean that the total energy is almost double?

That's how we treated it in earlier parts: Edrag E lift and total E is the sum.

Does this mean that we should therefore divide the ratio of E_plane / E_car by two since only half the energy is consumed by fighting drag?

I think we multiply it by two, because the ratio is comparing total energies, and we only considered the energy of drag, when infact the total energy a plane expends is twice that.

A curious thought I had after reading this was "what is the fuel efficiency of a flying car?"

Probably really poor. Cars are designed not to fly unfortunately, so it would take a lot of energy to make a car fly, if it's even possible at all.

Cars are designed to create downward force.... if you could fly it people would start to take off at high speeds.... and that wouldn't be so good.

Actually that's why drag racers are shaped the way that they are....so that when traveling down the track they don't take off and lose control.

$$A_{\text{plane}} \sim 6\,\text{m} \times 6\,\text{m} = 36\,\text{m}^2,$$

whereas

$$A_{\text{car}} \sim 2\,\text{m} \times 1.5\,\text{m} = 3\,\text{m}^2,$$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{\text{drag}} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

So since a plane is more streamline and this effect dominates the other errors it would be more efficient than a car? I had kind of understood that planes were much more efficient than cars when you take into account the # of people on the plane. I was on a flight in the last couple of years where the pilot said we were using about $20 in fuel per person. That's only about 7 gallons of gas or about 200 miles on the highway

interesting, thanks!

I feel like a plane is much more efficient than a car, if you factor in how much quicker and farther they go, along with how much more they transport.

To Sat 6:23: I agree that a plane has other advantages such as speed or maximum distance, but I don't think these can be considered as components of the fuel efficiency that we're estimating here. I wonder though what kind of expression we could use to include some of these factors, at least to consider something like: Energy Consumed per Person vs Velocity.

I understand this is a class about estimation, and this information doesn't seem important. But out of curiosity, how are we able to determine when certain details (such as comparison of energy and drag) are negligible.

I wouldn't think the fact that the planes move quicker would mean they are more fuel efficient.

I agree, also just as a matter of the way to the two move. All things considered, the drag of th eplane from the air is going to be less than that of the air AND friction from ground on the car to start, not ot mention the fact that planes are more aerodynamic in design.

Needs capitalization.

I like "error" here better than "mistake" earlier. You can make errors in estimation, but I feel like a mistake is hard to come by since most of us are just guessing anyway!

I think this recap is really good to go over the parts we neglected in the calculation.

I agree that the recap is good, it does come across as a whole bunch of stuff going back and forth on sources of error and their mitigation.

$$A_{\text{plane}} \sim 6\,\text{m} \times 6\,\text{m} = 36\,\text{m}^2,$$

whereas

$$A_{\text{car}} \sim 2\,\text{m} \times 1.5\,\text{m} = 3\,\text{m}^2,$$

so the third factor contributes a factor of 12. The fourth factor contributes unity, since we are analyzing the plane and car making the same trip (New York to Los Angeles, say).

The result of the four factors is

$$\frac{E_{\text{plane}}}{E_{\text{car}}} \sim \frac{1}{3} \times 100 \times 12 \sim 400.$$

A plane looks incredibly inefficient. But I neglected the number of people. A jumbo jet takes carries 400 people; a typical car, at least in California, carries one person. So the plane and car come out equal!

This analysis leaves out many effects. First, jet fuel is used to generate lift as well as to fight drag. However, as a later analysis will show, the energy consumed in generating lift is comparable to the energy consumed in fighting drag. Second, a plane is more streamlined than a car. Therefore the missing constant in the drag force $F_{\text{drag}} \sim \rho v^2 A$ is smaller for a plane than for a car. our crude analysis of drag has not included this effect. Fortunately this error compensates, or perhaps overcompensates, for the error in neglecting lift.

I think knowing when to compensate for error is part of the art in the art of approximation. I think i tend to neglect a lot of things when estimating, but I guess the trick is to know when a value is rounded by too much or a guess was too wild.

I agree, I think a lot of times the hardest part of these problems is realizing when two errors will cancel each other out, is there any good way to recognize this?

And even if things don't cancel out, you have valuable info if you know your answer is a little big or small

This is a good question! I feel like if I were estimating and got an answer, I wouldn't be able to tell if I made any errors (or where I made them). It would just be luck that they canceled out.

I guess my question would be...how much can I overcompensate? What's reasonable? a factor of 10? and if i overestimate should i just assume that something else must be underestimated? or that I can neglect another variable?

I think it is just a fun fact of estimating. It's pretty easy to change the estimate later by finding values or getting better estimates, but I don't think it's that hard to realize something is above or below an estimate (like mph of a car at 50 is an underestimate or something).

Still, how are we to tell if our overall estimation is over or under? Are we to tell somehow based on the overall direction of our estimation?

I wonder how the friction of the wheels on their chassis adds to the energy loss equation along with drag.

I think you should add to the reading what we talked about in class about how planes are not a direct substitute for cars. Like you said, people from the east coast wouldn't go to California as much if they had to drive. Therefore, there's something to be said about plane's having a larger affect on the environment due to their generation of sometimes unnecessary travel.

I feel like a lot of these estimations are still based on numbers I wouldn't know off the top of my head, which also help guide whether you round up or down to cancel errors

5
Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. “When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?” asks Laura Morosini.

Before continuing, explore the following question:

▶ *What is the most egregious fault in the comparison between Exxon and Nigeria?*

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria’s GDP (assuming the flow remains

I feel like this is just repeating already known information. If I hadn’t have known this, I would have been screwed in 8.01

> When you say "this", are you referring to the reading or to particular parts of it? I don’t claim that the ideas in this section are previously unknown, only that most people don’t use them. Hence the need for the section.

Dimensionless are the best. When there are many different terms to rectify with the numerator and denominator, it gets a little tricky.

Is it still this high? I think 2006 was at the height of the peak oil scare, when oil was well over $100/barrel...

5 Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

▶ *What is the most egregious fault in the comparison between Exxon and Nigeria?*

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

We're starting a new unit. Read the first two sections of this chapter for Tuesday's memo.

I hope that you don't find too much to fix in Section 5.1 because it's borrowed from _Street-Fighting Mathematics_ (MIT Press, 2010), which came out in print last week. But comment freely about it (and the other section) and let the chips fall where they may.

I like how this example/topic isn't about formulas or science, it makes for a nice change of pace and adds variety to the book.

I was going to cheer also, but then I noticed that we'll be back to it in the next few sections. But it is a nice change.

I like that we're relating to something I know better.

I'm an engineer and I'm getting kind of queezy because I didn't do all that well in 14.01 :-/

Looking at the section titles, I recognize almost all of them from Street-Fighting Mathematics. Have they been changed at all?

Why is this bracket here?

Probably a placeholder that got left?

Or maybe it's for a citation for the next passage

If this section was taken from another book, then the TeX could reference a citation that does not exist in the references/bibliography for this file.

5
Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

▶ *What is the most egregious fault in the comparison between Exxon and Nigeria?*

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

I've never had an econ class in my life, but isn't GDP yearly while net worth is the total amount the company would be worth if you sold it, not how much it makes in a year?

Yes, you are correct about both. Net worth is the value of a company (assets - liabilities) if you sold it right then.

Regarding, "I've never had an econ class in my life," maybe that is an advantage. When I gave a workshop one time and used this example to illustrate the importance of dimensions, many participants (who are faculty at leading research institutions) insisted that there was no problem!

Does this mean Exxon in Nigeria, or the worldwide corporation? I'm not sure how these things work in specific countries.

Worldwide. The comparison is trying to show how poor a country is in comparison to a single company (globally)

This is a poor comparison. net worth is at a particular point in time while GDP is cumulative over an entire year. It might be better to compare Exxon's annual revenue to Nigeria's GDP

I agree but the comparison isn't too ridiculous. It doesn't compare two similar things but, if anything, it puts into perspective the size of Exxon.

Fun fact: Exxon Mobil is the world's second largest company by market capitalization (that's share price*number of shares). It's also the most heavily weighted company in the Dow Jones Industrial Average, which is one of the most popular benchmarks of economic performance.

This is really interesting...and I think it emphasizes the idea of this quote and works as a great lead-in to the rest of the section.

Which begs the question, at least from me, what is the largest company?

I don't think I've ever heard a comparison like this. I usually hear GDP's compared to annual revenue or profits, which have the correct units.

5
Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

▶ *What is the most egregious fault in the comparison between Exxon and Nigeria?*

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

This is an interesting way to open the chapter. It made me want to read more to see how you'd tie in this excerpt to dimensions.

Yeah I really like this - it suddenly makes me very interested.

Agreed. It's a very general, non-technical overview that everyone can understand and ultimately relate to.

It's also, as someone mentioned earlier, a nice change of pace and topic. Airplanes are fun and good, but sometimes you can focus too much in one topic!

I also think the choice of topic (discussing the ethics or fairness) of globalization which is a very controversial topic further attracts the reader's interest.

I look at this somewhat differently - it seems all along we've been examining a few different examples. Although many were more physics/science based, we still did more than just airplanes.

This reminds me of District 9, and MNU operating in South Africa.

Yeah I am a huge fan of this introduction. It's simple to understand and straight to the point. The non Science thing is a good point too.

This is very similar to when people make claims like Harvard is worth more than small countries because of their $28B endowment. They don't mention that it took 400 years to build up and that they only spend 5% of it. It drives me crazy when journalists do this.

This is a great point that gave me a different perspective. Though, at the same time, you shouldn' compare Exxon to Harvard either. Exxon is driven so much more on a profit motive.

No matter what happens to the textbook, make sure to leave this part in!! This is the kind of stuff you always see on the news that is just nonsense. This is one parts of the class that isn't just about approximatoin, but about general lessons for life.

5
Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

▶ *What is the most egregious fault in the comparison between Exxon and Nigeria?*

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

This suggests Exxon's headquarters are in Nigeria- is this true?

i doubt their headquarters are in nigeria, but i'm sure they drill from there, so they "operate" there.

I don't think this suggests anything about Exxon's HQ (which isn't mentioned at all?)

My first thought was also that Exxon's base was in Nigeria, perhaps it could be changed to something like "The net worth of Exxon, who operates in Nigeria,..."

Changing 'the country' to 'a country' should remove most of that confusion, I think.

One of different units that might not really mean anything?

Haha, you catch on quickly.

Didn't she respond to your correction of this argument saying you were right but they were keeping it the wrong way to prove their point?

Oops, you mentioned that down below.

I'm unfamiliar with this name, is she a well known critic of globalization?

Other than this quote from "Impunity for Multinationals", I don't think she's a well-known critic. A quick google only came up with the one article.

I think it would be useful to add in parenthesis after this quote who she is.

I don't think a well known, educated critic would make this argument. It is so misguided.

what do you mean by this?

I think he's just asking what is the most obvious mistake in the argument/question posed by Morosini, not in her opinion, but in the way she argues it.

Sounds right!

5 Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

What is the most egregious fault in the comparison between Exxon and Nigeria?

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

I think the problem is Exxon operates in many countries. And perhaps serves more people than the population in Nigeria or serves more needs. On the topic of dimensions, I definitely wonder if the dimensions are the same since I am sure net worth and GDP are not measure in the same capacity

Right. I don't think the comparison of GDP to net worth makes sense. And isn't this chapter on dimensional analysis? I guess I'm cheating...

I agree as well, as GDP is the amount of money a country takes in per year while net worth is the amount a company is work after several years of worth. I wonder if the two are still comparable...

I dont think the two are comparable. I think GDP might possible comparable to net income in terms of the units but like mentioned earlier they represent two different things. That to compare net worth's of the two you would have to include all the natural resources, inventory, other available resources, etc. that Nigeria has.

I agree with the above - if you sold Exxon it would come with a lot more than its current bank account earnings for one year. If you sold Nigeria it would come with a lot more than its GDP.

Both are measured in currency, but I don't think the two can be compared easily.

this statement is a little confusing and not very helpful.

The field of economic reviews I think.

What field? The comparison?

I think the field of faults in the comparison

i dont even know what this means either.

maybe he's saying that: the economic field is competitive but in this particular case, the justification is wrong.

I took this sentence to mean that there are multiple errors in Morosini's argument/question, but we're going to focus on the most obvious one.

It seems like he's being sarcastic, saying there are so many errors to choose from.

5
Dimensions

5.1	Economics: The power of multinational corporations	85
5.2	Dimensionless groups	87
5.3	Hydrogen atom	91
5.4	Bending of light by gravity	101
5.5	Buckingham Pi theorem	107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

What is the most egregious fault in the comparison between Exxon and Nigeria?

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

GDP is more a measure of economic activity and production than monetary flow. It's specific to what was produced in a nation with regards to goods and services rendered. Not really an important thing, I know. I'm just nitpicking.

This is a great introduction to dimensional analysis, and I like how units play such an important role in this comparison

this is a little pedantic...

But it makes the point

And is interesting :)

But adds nothing to the discussion...

It's merely discussing the terms that are involved in the statement. Nothing "pedantic" about it. In fact, we regularly ask him in our comments to define stupid stuff that we could look on our own all the time.

That's interesting. I never considered how meaningless this measurement was.

This is a really interesting comment- it definitely made me think!

I agree. I really liked the way this was worded.

I like how this points out how arbitrary are units of measurement are

an interesting way to put it...

5
Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

What is the most egregious fault in the comparison between Exxon and Nigeria?

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

It doesn't seem entirely arbitrary.

Not entirely, but think about it this way: our life has nothing to do with how long it takes the earth to go around the sun, but one year is a crucial unit in our lives.

i don't think it's arbitrary at all. agriculture always has an effect of the economics of the country, and in some countries, it's a huge factor. agriculture depends upon the rotation of the sun.

I agree that a year isn't entirely arbitrary; every civilization has had some measure of time based on the seasons or the stars, which makes sense to me. I think the 365-day year was mostly to make bookkeeping easy for everyone (why have some years have two winters and others two summers?) I agree with what the reading says, however; to me it would make more sense to have a larger sample of, say, 5 astronomical years to smooth out the data.

I've thought about this on occasion, and how we celebrate our birthdays and say that we were born on some day, when the construct of a year has little effect on our modern lives. I don't think it is arbitrary in this sense though: a day as 24 hours seems not arbitrary because of our sleep cycle, it is so intrinsically tied to our bodies. If we say the measurement of day is arbitrary, then there is not much left to discuss, as anything else can be deemed arbitary. A year's direct influence on our body is so much more less drastic, but it does affect the changing of seasons, which directly affects agriculture, which, for much of history until now, has had a central impact on life.

I think this is kind of unnecessary.

yeah I don't really think this is necessary...maybe the same thing could be stated in a shorter way?

This is a great way to get people to start thinking about dimensional analysis. if we start looking at a decade, then the GDP all of a sudden looks so much larger

And sets up why the direct comparison doesn't work as well

I think using decades is common.. but I also wonder if that is a good idea for modern economics when things change so fast every year, if not every month. Is it OK to generalize it that much.

5 Dimensions

5.1 Economics: The power of multinational corporations 85
5.2 Dimensionless groups 87
5.3 Hydrogen atom 91
5.4 Bending of light by gravity 101
5.5 Buckingham Pi theorem 107

5.1 Economics: The power of multinational corporations

Critics of globalization often make the following comparison [] to prove the excessive power of multinational corporations:

> In Nigeria, a relatively economically strong country, the GDP [gross domestic product] is $99 billion. The net worth of Exxon is $119 billion. "When multinationals have a net worth higher than the GDP of the country in which they operate, what kind of power relationship are we talking about?" asks Laura Morosini.

Before continuing, explore the following question:

What is the most egregious fault in the comparison between Exxon and Nigeria?

The field is competitive, but one fault stands out. It becomes evident after unpacking the meaning of GDP. A GDP of $99 billion is shorthand for a monetary flow of $99 billion per year. A year, which is the time for the earth to travel around the sun, is an astronomical phenomenon that has been arbitrarily chosen for measuring a social phenomenon—namely, monetary flow.

Suppose instead that economists had chosen the decade as the unit of time for measuring GDP. Then Nigeria's GDP (assuming the flow remains

I guess now it makes sense why the sentence about the year being the unit of time chosen to measure social phenomenon was in the previous paragraph...it gives a nice transition to this paragraph.

I think in the first paragraph it should be stated that GDP is a flow (which is does... good!), but additionally, I'd like to to make clear that Net Worth is an amount. I think it shouldn't be stated much later- once it is mentioned, it becomes clear the comparison and everything flows a bit better.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

Comments on page 2

Wait, but that means you're comparing apples to oranges now: Nigeria's GDP for a decade, and Exxon's for a year. Wouldn't you have to compare it to the monetary flow of Exxon in the past decade...?

Ha, nevermind, didn't realize that the original quote compared GDP to net worth. That was silly.

I feel like I made the same observance. I was thinking that these two things are incomparable..but then later in the reading it explains what I was thinking.

but these can't be compared, one is over a period ten times the other

I don't get this reported business. Under what circumstances would the reported number be different than the actual number?

I think he means reported as just "1 trillion" the number without any time scale attached. Because without the time scale it's easy to lose sight of the fact that a net worth doesn't depend on a unit of time and a GDP does.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\text{m}\,\text{s}^{-1}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

This is an interesting point.

I missed this in my first read through of the excerpt above, and it pushes a rather strong argument against the "critics".

That is, in this case, Nigeria "makes" as much in a year than the entire value of a large corporation...

I agree. I hear some educated people make the same sort of error when they're arguing for/against certain policies. I wonder if there's some psychological block that causes us to make these sort of errors.

This intro was especially noteworthy for me because I have heard many times how a corporation has a higher net worth than a country's GDP, but never had I stopped to realize what that statement incorrectly indicates.

As for the "psychological block" you alluded to, maybe those statements are trying to say that a corporation is worth so much that it could run an entire country for a year?

Indeed, despite my economics nackground, I had not considered the arbitrarity of our definitions and the year.

I don't have an extensive background of economics outside of 14.01 and 14.02 but this example really drove home how important it is to really think about what we are comparing when we use invented forms of measurement.

I think my favorite part about this is the tone it seems to convey with towering over puny assets - it gives almost a comical view at how easy it is to play with units to change a number presented with improper dimensions

I think someone brought this out above, but at this point still, does this contrivance that Nigeria "makes" more even really have a point? Because we are still comparing GDP(a rate), to net worth at a time. Shouldn't we compare GDP to GDP, and net worth to net worth, altough, measuring the "net worth" of Nigeria is a complex equation?

EDIT: oh nvm, the paragraph right below talks about this, haha.

Why are we comparing a net flow to an absolute amount? It doesn't seem like we want to compare them at a moment in time.

wouldnt Exxon's gdp change also, if we are choosing different time specs? or are you implying the 2 are measured in differetn ways? in that case, how is exxon's gdp measured?

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

this makes me wonder what would be an appropriate time frame (unit). in fact, is there an "appropriate" unit, people can play around with the units and therefore convey a totally different message

That's still pretty big!!!

Interesting. this is a good cautionary tale of looking beneath the words/numbers in statistics. Care always needs to be taken to reacting to these kinds of figures. I anticipate this unit will make us a much wiser consumer of information.

I agree - i really like this

You should read the book "How to Lie with Statistics" by Darrell Huff. Its about all the tricks that can be pulled with wording and numbers in statistics if you aren't careful about what you're reading.

I agree, this was a very insightful and interesting example of the concept of dimensions.

This is a wonderful argument and explanation.

It would be interesting to see how long it took Exxon to attain that net worth- i.e., did Exxon build itself up in a year or in 20 years?

Thats a good point, and what did their growth look like? similar to that of a nation?

Does this mean you can compare the GDP of Nigeria over the time span that Exxaon has been making profit and it be a valid comparrisson?

wouldn't the ratio of nigeria to exxon GDP scale relative to whatever time you are considering?

Is this comparison even valid? Is there a point to comparing between something that changes with time and something which doesn't (unless Exxon bankrupts or something that dramatically changes their worth)?

No, the whole point here is that the comparison is totally invalid because the units are different.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

A good way to put it for people who don't do a lot of work with numbers.

Adjectives are here to save the day and spice up the writing.

Alright I'm guessing we're going to end up comparing annual revenue to GDP.... I wonder if you could get the net worth of Nigeria lol assets? cash flows? Is it possible while remaining politically correct? what about human capital? Is it fair to compare citizens to employees?

I don't follow this logic- if you use the same time period, the GDP can be compared

Again, a good point. Reminds me of examples when people use lightyears as a measure of time...

This is a very interesting discussion - great way to open and motivate the topic of dimensional analysis.

Very good point. Claims like the one you mentioned above are used every day, and it's really up to the individual to pick out what is a lie and what is a twist of truth and what actually is truth.

Ever since I can can remember, we have always been told to look at our units when comparing values. However, I find it extremely amusing how you show it is a very obvious error when you compare money, with money per time.

Er, this makes it sounds like the problem is the unit of time chosen (i.e., using an astronomical phenomenon), not just that the problem is comparing flow to net worth.

Maybe he is trying to point out the error in both. 1) we should probably use a time-scale that is characteristic of the phenomenon that we are trying to describe. Does anybody know what timescale this might be for GDP since the year is what is most often used in accounting? 2) One should not compare GDP and net worth because they are two different properties.

I agree with 5:58, the main point of this couple sentences is muddled by the phrase about astronomical phenomemon. It's ok that we use the year to measure GDP and we could also use something like 5-year periods to measure it. (I wouldn't really consider a decade a different "astronomical phenomenon" from a year since it relies on the same happening.) It's that we're comparing a rate to a time-independent quantity that's the problem. Mixing things that are both rates but with different time periods would be a different problem.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

Wow. I was totally fooled by this argument at first. I didn't even think to compare the units the two numbers were in. But this is a good point on how people make statistics mean whatever they want.

i didn't get this at the start, but it is the most important part. why can't the chapter start out with this?

There's a difficult tension between writing a handbook and a textbook. The handbook is for those who already know the material and want a quick memory jog. The textbook is for those learning it for the first time.

The two are quite different, because just telling people something is not a good way for them to learn and understand it. It's much better to build a structure in which the ideas come into play, and the reader forms their own understanding (with guidance).

But once they form the understanding, they are now closer to the experts, and would rather have the core ideas presented right away – but it might not have been instructive if it had been done that way.

Perhaps with the web there is a way to merge both kinds of books into one document?

This made me re-look at the quote in the beginning...i hate that people are able to get away with this kind of lie, arg.

Given that most of our examples have been scientific so far, and so few (if any) have touched on the social sciences, this is a nice change of pace.

I completely agree...its definitely relating approximation to everyday world. This was a useful example and definitely examines the usefulness of approximations.

I couldn't find where to put this comment, so I'll agree to this point (I do like economics examples), and also add that i think this first example was really useful in starting this new chapter.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

Financial analysts often use income data for several years to project a company's lifetime net worth. Perhaps this could be used on countries. This method is called the discounted cash flow analysis.

I don't think it's possible to calculate a realistic net worth of a country using a DCF. DCF itself for a company is very theoretical and sensitive to assumptions (about discount rate, projected cash flows, etc.), and I could only imagine this sensitivity increasing if looking at a country.

I agree. We are able to come up with projected cash flows from the current situation of a company and a market but it would be very difficult to look at the change in rate of a country. Being the second derivative implicitly makes it more sensitive.

It would be interesting here to find a comparable dimension, to show just how flawed it is (like, figure out the monetary flow rate of Exxon)

Or maybe the net worth of Nigeria. That seems like it would be a huge value, counting all sorts of infrastructure.

These ideas are intriguing... I would have never thought of this comparison without being enlightened by these facts.

Excellent way of transitioning from on opening introduction to the topic of this chapter!

I agree, I'm enjoying this example a lot.

That was a great example to lead into this paragraph. By showing the problems with comparing things with different units you show why it is important to take care when comparing data.

Yep, I also agree. Interesting, clear, and gives me reason to care about the topic at hand.

This sentence is a good example of the difference between units and dimensions.

I think this sentence should be in BOLD and called out instead of hidden in a parenthetical, since it's an important and maybe subtle point.

I think this should say "A dimension is a general..."

Whoops, never mind. My correction is completely off.

lol

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

can you look at a quantity and know instantly that it is either a dimension or a unit or it is also a factor of context?

You can tell right away. If you can sensibly write a number before it, then it's a unit; otherwise it's a dimension. For example, "5 meters" makes sense but "5 length" makes no sense.

this comment would be good for a side bar.

Good clarifying statement. I think people interchange these in everyday speech, so this sentence helps clears up some possible misunderstanding.

I also like how you define the topic very early in the reading, even as you it in an example.

yeah, this is a great point. You simply can't compare a flow or rate with an amount. This is done a lot, and if you're not careful, you might not be getting the facts straight.

While I'm glad that you point this out early on, I wish you would have expanded on it more. I'm not sure I entirely understand what you mean.

Rereading the above example of money versus dollars, I now understand what you mean.

i'm confused between a dimension and a unit.

I'm not sure but I think a dimension is a general characteristic like time, distance, etc. A unit is a measurement system for dimensions (i.e. inches, meters, cubits, etc, for measuring distance)

might be good to italicize dimension and unit here

So, what would be a more proper comparison? Annual Revenue of Exxon vs. GDP of Nigeria? I think a revenue comparison is more apt than a profits comparison

I actually learned this in one of my Economics class and it is really cool that this concept is mentioned in this class as well.

Never mind, I see

steady from year to year) would be roughly \$1 trillion per decade and be reported as \$1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes \$2 billion per week, reported as \$2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

> **Problem 5.1 Units or dimensions?**
> Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

so this is equivalent to comparing velocity and distance?

Never mind, the very next paragraph answers my question.

Answered previous question. I see where this paper is going now

Why is this bracket here?

This is also the second time in this reading that this has happened?

Maybe there is supposed to be some reference to another section/appendix/something? It seems odd that there are multiple instances of hanging brackets...

> My fault. I did not include the citation in the bibliography database, so the reference pointed nowhere. I'll fix that and the other spots (though I had done it before, but I must have missed a few).

What is a "category mistake" besides something that is "guaranteed to generate nonsense"

they belong to different categories, so it doesn't make sense to compare them. ie: apple and orange are 2 different things, and it doesn't really make sense to compare them

why category though? it's just wrong dimensions. hence, dimensional analysis.

I think i like dimension better than category as well

This clause seems unnecessary, why can't you just leave it at "Because their dimensions differ, this is guaranteed to generate nonsense"?

...or powerful propaganda for uneducated people

its true, a lot of platform statements are based on incomparable quantities or biased statistics

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 **Units or dimensions?**

Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk 1.5 m s^{-1}—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

i don't consider this nonsense. though the time is in some ways "arbitrary," it's still important to compare with that value of a "year."

But as we saw, it makes a big difference what unit of time we choose. If both sides of the comparison don't use the same unit of time then our answer doesn't make any sense.

I think about the amount of money something will cost me in comparison to my annual (or daily or monthly) salary all the time and it's not nonsensical. Similarly, house prices are compared to annual income, assets to annual income, PE ratios... It's not nonsense if we have a sense of what we're comparing. If we have a sense of a year, and a sense of the size of companies then it can be a useful comparison. It doesn't mean that the magnitudes have to be comparable (though that's probably the rhetorical point of the comparison and the danger of misused numbers).

I loved this example as the start of this section- it's really intriguing.

however, one intimates the other.

This is a nice problem to start off with to get acquainted with what we're doing next.

It's a bit sad that this actually took me a minute!

These are units

This question would be more effective as "indicate which of the following are units and which are dimensions: ..." The way you have it is rather leading.

These, and energy, are dimensions.

These are dimensions which are composed of the units: meters, kg, and seconds.

This paragraph comes out of nowhere...it's unnecessary...and i think that it's actually distracting to the discussion on dimensions because it's just thrown into the middle of a section on economics (which it has nothing to do with.)

it sounds so stupid when you say it like this

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

I really appreciate this comment because people say things like this all the time. I am quite positive no one really things about the importance of dimensions.

I was just going to comment that I've never heard someone say something like this. The GDP/net worth example is a good one and seems realistic, but I don't think I've heard anyone compare speed and distance incorrectly.

Another thing that gets me- people measuring distance in units of time... As in, "the grocery store is 10 minutes away." No, the grocery store is half a mile away and it will take 10 minutes to walk there. This happens ALL THE TIME

I feel like when someone says the grocery store is 10mins away the "...by X method of travel" part is implied, but I agree people mix up units a lot in general conversation although I haven't quite heard one like the example given here.

I think the grocery store thing might be a better example (despite what is obviously implied) because I've never heard anyone say anything like the example quotes about the Empire State Building. Or if I have heard something it is different enough that I read those quotes and was rather confused by how anyone could let that pass...

In the grocery story case, I feel it's because people don't actually care how far away it is in units of distance but rather how long it will take them to get there.

I agree about never hearing the example here before. Perhaps it would be more useful to put an example that would be more relevant/related to students like the grocery store example given above.

I have also never heard this anything like this before. It seems so unreasonable, I think that anyone intelligent enough to use Meters, would not make any mistake like that.

I agree that nobody says the speed versus length comparison. That's why I chose it: because it's so obviously bogus that no one would say it. Yet the equally bogus GDP versus net-worth comparison is made so often. Somehow I should make it clearer why I chose the speed versus length comparison.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

I was surprised that the speed vs length example was even in here; i read it and said, "There is no way anyone would actually say that. It would be ridiculous." The grocery store one would probably not be as good as an example because I feel it's semantics. I know, personally, I say "It's 5 minutes away", which is just me rephrasing "It will take us about 5 minutes to get there."

> The ridiculousness of it is exactly why it is in there. Nobody would say such a thing about speed vs length, but the equivalent comparison about GDP versus net worth is made all the time, and I want people to see, by analogy, how ridiculous that one is.

I was surprised that the speed vs length example was even in here; i read it and said, "There is no way anyone would actually say that. It would be ridiculous." The grocery store one would probably not be as good as an example because I feel it's semantics. I know, personally, I say "It's 5 minutes away", which is just me rephrasing "It will take us about 5 minutes to get there."

Even though it's based on certain implicitly agreed upon conventions, this immediately makes me think of how people say that places are "2 hours north" or that they're "ten minutes away from you." I realize that it's followed by an unspoken "at a constant speed near the highway limit" or "at my current walking pace," but it's still something that sticks out in my mind.

I've never heard something like this without time also mentioned. Relating velocity to distance through time.

I think this is a good example to show the flaw in the earlier comparison. I know I'm not too familiar with GDP and net worth, etc. so this example really solidifies the issue.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

why did you write the same example (essentially) twice?

I guess he try to pinpoint that if you use different units, ms-1 and m/hr you would get totally different conclusions, once is smaller than the height of the Empire State building while the other is much larger than the height. Yet the speed is same but just expressed in different units

(OP) this makes a lot of sense actually, thanks

It's the same structurally, but in the Exxon/Nigeria form it is easy to not see the problem, whereas in the speed versus length form it is much easier to see. So, I was trying to make the first example as obviously problematic as the second one.

I feel though you cannot compare Nigeria and Exxon, sure Exxon might operate in Nigeria, but it has plenty of offices in Houston and London that generate themselves millions of dollars in sales, of which Nigeria has no access to in these regions. So I'm not sure weather arguing for either case is correct

I don't think he's arguing for either side, rather pointing out the error in comparison.

I think the comparison is just to show that a company has a higher production than an entire country. But here he's just pointing out the error.

What author?

What!!!?? that's so ridicoulous! Why is that okay? Does the strength of a argument suddenly compensate for the accuracy?

I feel like the accuracy should largely determine the strength of the argument.

It compensates if your reader is uneducated or easily manipulated...

This is completely repulsive. Making a mistake is one thing, but keeping in a mistake because it might fool a few people into believing your point is horrendous. It's akin to making up experimental data to suit your hypothesis. Awful.

This is why the only articles I read now are science ones. I hate journalists.

steady from year to year) would be roughly $1 trillion per decade and be reported as $1 trillion. Now Nigeria towers over Exxon, whose puny assets are a mere one-tenth of Nigeria's GDP. To deduce the opposite conclusion, suppose the week were the unit of time for measuring GDP. Nigeria's GDP becomes $2 billion per week, reported as $2 billion. Now puny Nigeria stands helpless before the mighty Exxon, 50-fold larger than Nigeria.

A valid economic argument cannot reach a conclusion that depends on the astronomical phenomenon chosen to measure time. The mistake lies in comparing incomparable quantities. Net worth is an amount: It has dimensions of money and is typically measured in units of dollars. GDP, however, is a flow or rate: It has dimensions of money per time and typical units of dollars per year. (A dimension is general and independent of the system of measurement, whereas the unit is how that dimension is measured in a particular system.) Comparing net worth to GDP compares a monetary amount to a monetary flow. Because their dimensions differ, the comparison is a category mistake [] and is therefore guaranteed to generate nonsense.

Problem 5.1 Units or dimensions?
Are meters, kilograms, and seconds units or dimensions? What about energy, charge, power, and force?

A similarly flawed comparison is length per time (speed) versus length: "I walk $1.5\,\mathrm{m\,s^{-1}}$—much smaller than the Empire State building in New York, which is 300 m high." It is nonsense. To produce the opposite but still nonsense conclusion, measure time in hours: "I walk 5400 m/hr—much larger than the Empire State building, which is 300 m high."

I often see comparisons of corporate and national power similar to our Nigeria–Exxon example. I once wrote to one author explaining that I sympathized with his conclusion but that his argument contained a fatal dimensional mistake. He replied that I had made an interesting point but that the numerical comparison showing the country's weakness was stronger as he had written it, so he was leaving it unchanged!

A dimensionally valid comparison would compare like with like: either Nigeria's GDP with Exxon's revenues, or Exxon's net worth with Nigeria's net worth. Because net worths of countries are not often tabulated, whereas corporate revenues are widely available, try comparing Exxon's

Hmmm interesting. Now that you've exposed the Nigeria-Exxon article, perhaps the readers of your book can write letters en masse!

Only if Sanjoy is willing to name names...

Hardly sur[rising, this is what everyone does when posing an argument. Sad, but true. I too would love to see this article though.

It's entirely true people manipulate/present information in a certain way all the time to drive a point home. The interesting this about this example is that I don't think most people would even realize the error in the comparison to begin with.

Yeah, this example reminds me of how polling can be manipulated so that the poller's hypothesis can be confirmed. This is especially apparent in political polls when a specific party leader will conduct polls and claim that "80% of Americans agree with X,Y, and Z" but doesn't mention that the only Americans polled were members of his party.

wow. I really want to be surprised by this comment, but I just cant.

Sound Logic is not as powerful as clever rhetoric if you are trying to sway the masses

I really don't find this surprising. He doesn't really sound like he's from a technical background, which isn't surprising since he's making politically motivated comparisons.

Agreed -politicians only have to say things the majority of the voters will believe.... unfortunately most of them aren't talking dimension.

I thought so too. I wonder what he numbers are like. Probably a lot weaker since the author didn't use them in the first place.

I also thought this. I guessed and commented on these comparisons earlier in the reading although I didn't know they would be explained later.

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

Comments on page 3

Did you make the author aware of this?

And oil companies say they aren't ripping us off..

In what year was it 99 billion?

The quote that says 99 billion comes from an article that was published in 2002, if that's any help.

this seems like a silly statement to me: one is wrong, the other is right. of course it's better.

wow, I'm sure that author would be glad to hear that!

And I'm sure he also feels silly for using a flawed comparison instead of a valid one when both confirm his hypothesis!

It shows the power and how easy it is the manipulate the written word. Hard facts always wins out in the end.

when you wrote the author did you give him these numbers? or just say that the argument contained serious mistakes?

wow I'm really surprised. I thought the author was making the "mistake" on purpose. I guess their are other flaws than time when comparing assets and annual cash flows. But really what is a country and what is a company? Maybe there are even more errors in making this comparison than in the units.

True, it may in fact be that exxon mobil has more political might than Nigeria

I agree but wasn't that the point of the statement in the beginning? I wonder why the author didn't just put this comparison originally.

You should send this to the author!

or just resubmit the paper with the correct calculations. seeing as teh author didn't seem to want to do it.

Did you explain that to the author?

He probably didn't listen even if it was tried. people are stubborn

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

Is this referring to the ignorant author referenced above? Because it seems like s/he was a totally unrelated anecdote to our nigeria example, but if they are one in the same I'd like to know as a reader.

I think after this you need to briefly explain what this - now valid - comparison means (tying up the "what kind of power relationship" quote).

This sentence sounds awkward. It took me a couple reads to understand what you meant.

suggestion: The condition that compared quantities must have identical dimensions is necessary but not sufficient"

On this point, I would think there's a better statistic than revenue to compare to GDP.

I understand your point, that we need to mind our units, but it doesn't mean that we can't compare quantities with the same dimension but different units. For instance 1 meter > 3 ft, etc.

Earlier in this sentence, he says, "the compared quantities must have identical dimensions", so I think that covers the case where you're discussing length with both meters and feet.

Yes, but my point is its possible to make valid comparison without matching units (and therefore, that matching dimensions is sufficient). It's certainly easier with matching units, but not necessary to have them.

saying that 1m/3ft is not ignoring units! in trying to figure out which one is larger...you do the conversion in your head - you are just talking about unit conversions. And people should still be careful with this - remember the Martian lander than landed below the surface of the planet...

hahahaha i should have read just one sentence later...

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

Several similar cases have come up in recent history. Why don't people learn to keep units consistent? Learn to include units when giving calculations to others!!

Part of the problem also stems from having a globalized world without a common unit system. Same thing with language.

The situation is not great in America, but it's even worse in England (and maybe the rest of Europe). There, because the metric system is pretty much standard, all physics is done completely in the metric system and with SI units (meters, kilograms, seconds). So, teachers in high school tell students not to include units in the intermediate stages of a calculation but only put them in at the end.

This unwise advice used to drive me crazy, and I shared the Mars Climate Orbiter story very often with the students (one of the students had told me about it when he saw it reported in the newspaper).

In America, perhaps because feet and slugs and kilograms and meters are used in engineering, people are more careful to write units.

Was this the example that was mentioned in a couple weeks ago? I always like when examples in lecture correspond to the readings!

That was the example I mentioned in lecture.

I think this is an overused example...but maybe it's just because we go to MIT that we've heard it so often?

even so...its a good thing to keep in mind, so we never forget, just like the Tacoma-Narrows bridge (I see that video in every class too)

I have my doubts about whether this one was true. Don't they use simulations to test whether the program works? How could a gross mismatch in units slip through the simulations.

yeah I've also heard it said that they had to "type" the numbers into the program. who does that? they would have been read in automatically

I've never heard this before, so I'm glad for the example.

yeah I had never heard of this example either so it was really interesting and surprising when I read this

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

There's a Mishap Investigation Board? What types of incidents do they investigate besides NASA?

That's a great name. It's an ad hoc thing that NASA does when things go wrong and a common military term, but we could probably use more investigations of mishaps.

I guess these brackets are for citations that haven't been filled in yet?

Good observation.

They are that. Though I thought I had found all those spots and had augmented the bibliography table – evidently not. That reference should be to, "Mars Climate Orbiter Mishap Investigation Board", Phase I Report, NASA (November 1999).

Another good example of metric-imperial confusion:

http://en.wikipedia.org/wiki/Gimli_Glider

I find it amazing that such simple mistakes are made on multi-million (or billion) dollar projects. Obviously, none of us will make these mistakes after taking this class...

It takes 2 to make these mistakes...always check not just your work but others

That page has double value. I was thinking about using it as part of a discussion on gliding. Life is short and we didn't get to it, but I'll mention it in the book, perhaps in both places.

How embarrassing.

This happened with some plane company as well. This kind of stuff should never happen.

Just shows that the little mistakes we make on psets and stuff can be made by people in real life too, and with much more serious repercussions.

Yeah I've heard about this before. Ever since reading this I've made it a habit to include dimensions whenever necessary.

That guy's life was ruined

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

Why doesn't the world just pick one and stick with it? Wasn't there a country somewhere that drove on the English side of the road forever and ever until like a decade ago one day the government said "switch!" (with a clever marketing campaign and many signs and patrolmen) and then they did? I might be making this up.

Wouldn't this cause even more of said problems? We're more or less in too deep in this situation and I think generally there aren't enough advocates or good reasons that outweigh the difficulties in the US to make this drastic change.

Possibly for a similar reason that everyone uses a different currency...

I think that was Sweden in 1967. The rest of Europe (except England, which was isolated by the Channel) used right-hand drive, so it made sense to do.

As for why left-hand drive started out, I heard various explanations while I lived in England. The most convincing, even if it isn't true, is that it originated on the original "highways" – roads for people to walk on. There, people walked on the left side and carried a sword (to defend against robbers or "highwaymen" as they were also called). The sword was normally on the left side, so that a right-handed person could reach across and pull it from the scabbard, and then would have the sword ready to protect their right side. Therefore, walk on the left to shield yourself from the rest of the road.

This choice, allegedly, was adopted when cars came to drive on the highways.

I love the examples used to reinforce this lesson!

Could this sentence be fleshed out to "Make sure you mind your dimensions and units...." and something about the cost of the MCO?

I think most people would have an appreciation for how much the MCO accident must have cost, so repeating it here might appear redundant. However, I sorta expected "and units" to be emphasized somehow (italics, bold, etc)

I like the succinctness of this sentence because it gets the point across well.

I agree, it acts almost like a concluding sentence. Correct dimensions and units are both necessary in approximation problems.

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons

Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

Annual budgets of cities vs. endowment values for universities. This is a constant source of town-gown tension and newspapers don't help the situation by trumpeting endowment values. Budgets would be a more relevant comparison.

A common reoccurring example is the recycling issue. I always think if the recycling numbers are true when compared to the numbers that produces the bottles themselves. Or the numbers of reusable aluminam cans, the production and transportation of them might as well overpower the recycling and production fo the plastic ones.

"now with 20% less plastic!"

Miles per gallon. When comparing the miles per gallon between a porsche and a pontiac, I've seen them say that they both can get 25 mpg. They fail to mention, however, that they were using premium gas for the porsche and regular unleaded for the pontiac, which I'm sure makes a big difference to someone who is looking to find the most fuelcost-efficient vehicle...

though i agree that this is an example of deceit, but i don't think the problem here is with units.

plus, it might as well be accurate. afterall, the person with the porsche is probably going to be the one that buys premium gas and the one with the pontiac will likely stick with regular.

This might not be relevant, but I heard that when they calculated the iron in spinach for the old Popeye cartoon they overestimated by a factor of 10.

Hehe

Did they overestimate the additional strength of his biceps due to the iron?

I find that this section was a really nice read and helpful to understand what we are going to be learning about.

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

In 2.006, one of the early topics was dimensional analysis and it was all about forming dimensionless groups. I thought it was really cool how we could understand literally none of the physics about a system but using the dimensions of its critical components formulate important dimensionless ratios to solve the problem. I'm looking forward to the angle this class takes with this topic.

Actually, I'm beginning to notice that a lot of my engineering classes actually use a lot of the topics in this class. In my power electronics class they mentioned symmetry a bunch of times. And I just realized that in every one of my classes, a dimensionless constant is always used somewhere.

Isn't this another name for scaling.

You mean when comparing two numbers of the same dimensions?

This is a really useful section, and this is a very useful tool. I think that this section should possibly be expanded with a few more short examples; I believe that learning this well can get anyone through any type of estimation; find the groups and find an answer. (It's also a way to find equations; you have these knowns and you need length: so we need to end up with a group that equates to L).

It actually took me some rereading to realize that this was referring to the Exxon example, since 'oil' wasn't specifically mentioned in that context (I was trying to think back to some estimation we had done on oil imports..). Maybe just stick with 'Nigeria' or 'Exxon example'?

Same here. Thanks for this comment!

Yeah I would consider rewording this.

It might be worth to at least mention, if not go into detail, the Buckingham Pi Theorem. For those who are interested in it here is the wikipedia page: http://en.wikipedia.org/wiki/Buckingham_%

Dimensionless ratio means all the units cancel right?

So it is more efficient to use ratios of quantities than absolute quantities? that may be help for for quick calculations but isn't that less practical?

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

Well, you mean for estimation purposes, right? Not all the time...

I mean all the time! The universe doesn't care about our system of units (e.g. whether we measure speed in meters/second or furlongs/fortnight), so our descriptions of the universe shouldn't care about the system of units either. And the best way to ensure that our description doesn't care about the system of units is to use dimensionless quantities, for they are invariant to a change of units.

I was starting to see this pattern in proportional reasoning, and I'm glad we're looking at it now

I guess this is true actually–if you think about it, if someone told you a plane travels 500mph, you wouldn't instantly be like oh 500mph i know exactly how fast that is. You'd think to yourself, ok a car goes around 65 mph on a highway, then think about the ratio of speeds to get an idea of how fast that plane goes.

I think this is a great explanation! I think the paragraph could benefit from a narrative comparison like this.

I'm really glad that someone in class the other day asked about dimensions and proportional reasoning and why they don't always match in that case... I think you should include something about that in the reading (it also helped clarify the and proportional symbols)

Yeah, this was a huge help. And the quick example presented a couple of comments up helps to make this link between proportional reasoning and dimensionless ratios.

Yeah, because without the comment stated above with the example I would have been kinda confused as to why you don't care about dimensions. So by dimensionless quantities we are talking about dimensionless ratios? Not just dimensionless things in general like comparing degrees to a constant.

On that note, I saw the class notes from one of the other classes taught by Sanjoy (Streetfighting Mathematics maybe?) and there was a good little table explaining the difference between and = and stuff. I think that would be good to include somewhere.

That's a nice explanation.

annual revenues with Nigeria's GDP. By 2006, Exxon had become Exxon Mobil with annual revenues of roughly $350 billion—almost twice Nigeria's 2006 GDP of $200 billion. This valid comparison is stronger than the flawed one, so retaining the flawed comparison was not even expedient!

That compared quantities must have identical dimensions is a necessary condition for making valid comparisons, but it is not sufficient. A costly illustration is the 1999 Mars Climate Orbiter (MCO), which crashed into the surface of Mars rather than slipping into orbit around it. The cause, according to the Mishap Investigation Board (MIB), was a mismatch between English and metric units [, p. 6]:

> The MCO MIB has determined that the root cause for the loss of the MCO spacecraft was the failure to use metric units in the coding of a ground software file, Small Forces, used in trajectory models. Specifically, thruster performance data in English units instead of metric units was used in the software application code titled SM_FORCES (small forces). A file called Angular Momentum Desaturation (AMD) contained the output data from the SM_FORCES software. The data in the AMD file was required to be in metric units per existing software interface documentation, and the trajectory modelers assumed the data was provided in metric units per the requirements.

Make sure to mind your dimensions and units.

Problem 5.2 Finding bad comparisons
Look for everyday comparisons—for example, on the news, in the newspaper, or on the Internet—that are dimensionally faulty.

5.2 Dimensionless groups

Dimensionless ratios are useful. For example, in the oil example, the ratio of the two quantities has dimensions; in that case, the dimensions of the ratio are time (or one over time). If the authors of the article had used a dimensionless ratio, they might have made a valid comparison.

This section explains why dimensionless ratios are the only quantities that you need to think about; in other words, that there is no need to think about quantities with dimensions.

To see why, take a concrete example: computing the energy E to produce lift as a function of distance traveled s, plane speed v, air density ρ, wingspan L, plane mass m, and strength of gravity g. Any true statement about these variables looks like

this seems like an extreme (not to mention counter-intuitive) statement

this seems like an extreme (not to mention counter-intuitive) statement

i agree. often units can help guide you towards the right answer and are easier to deal with than dimensionless ones.

Its about getting things from one place to another using only unit. It wont always work, but I can only imagine that its very useful

What a fascinating statement

It would be nice to apply the same analysis for SHM to this problem and get a crude expression for the energy.

dont really understand what a "true statement" do you mean an equation that is valid

I think 'Any true statement' is hyperbolic. Do you mean "We could write a true statement about these variables that looks like..."

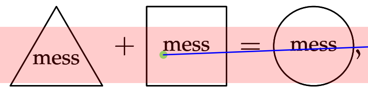

where the various messes mean 'a horrible combination of E, s, v, ρ, L, and m.

As horrible as that true statement is, it permits the following rewriting: Divide each term by the first one (the triangle). Then

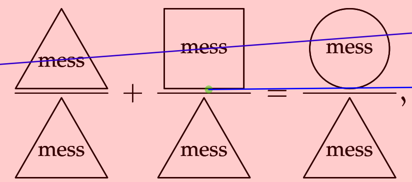

The first ratio is 1, which has no dimensions. Without knowing the individual messes, we don't know the second ratio; but it has no dimensions because it is being added to the first ratio. Similarly, the third ratio, which is on the right side, also has no dimensions.

So the rewritten expression is dimensionless. Nothing in the rewriting depended on the particular form of the true statement, except that each term has the same dimensions.

Therefore, *any true statement can be rewritten in dimensionless form.*

Dimensionless forms are made from dimensionless ratios, so all you need are dimensionless ratios, and you can do all your thinking with them. Here is a familiar example to show how this change simplifies your thinking. This example uses familiar physics so that you can concentrate on the new idea of dimensionless ratios.

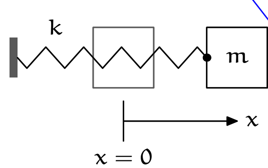

The problem is to find the period of an oscillating spring–mass system given an initial displacement x_0, then allowed to oscillate freely. The relevant variables that determine the period T are mass m, spring constant k, and amplitude x_0. Those three variables completely describe the system, so any true statement about period needs only those variables.

haha this is exactly how physics works in my head

This is how most of my class work looks like.

I don't know if you are missing a quote at the end or this is just a typo

mm I read it as "prime a" but I guess that doesn't make sense... lol

How is this useful?

I'm not sure using shapes here is the best way to communicate your point. I understand what you're saying, but I can also see how other readers (perhaps non-MIT students) might be confused. I would consider reworking this.

Or, you could keep the shapes and just make note that same shapes imply that the two "messes" have identical dimensions.

I actually think this is really clever and makes the point clear.

I also like the pictures, it reminds me of high school when we learned about how to convert dimensions by "canceling them out." The point was that it didn't matter what you were multiplying (3 squares divided by 7 circles) as long as the top equaled the bottom.

I had to look at it twice to get it, but I really like visual examples and this one make sense to me

The shapes serve as a good visual, but I think a quick, basic example should be included as well (much more simpler than the spring example).

I think it was a great example, especially for the visual learners like myself that did not benefit as much with all of the math equations and essays from previous chapters. Also, it requires no context or problem set-up.

So what kind of dimensionless constant could you use for comparing the company and the country?

This explanation is still a little fuzzy to me... I don't really understand how we can go from 3 "messes" to 1 dimensionless statement.

How do we know this is true?

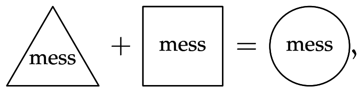

where the various messes mean 'a horrible combination of E, s, ν, ρ, L, and m.

As horrible as that true statement is, it permits the following rewriting: Divide each term by the first one (the triangle). Then

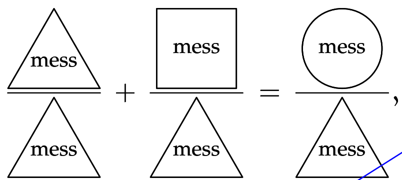

The first ratio is 1, which has no dimensions. Without knowing the individual messes, we don't know the second ratio; but it has no dimensions because it is being added to the first ratio. Similarly, the third ratio, which is on the right side, also has no dimensions.

So the rewritten expression is dimensionless. Nothing in the rewriting depended on the particular form of the true statement, except that each term has the same dimensions.

Therefore, *any true statement can be rewritten in dimensionless form.*

Dimensionless forms are made from dimensionless ratios, so all you need are dimensionless ratios, and you can do all your thinking with them. Here is a familiar example to show how this change simplifies your thinking. This example uses familiar physics so that you can concentrate on the new idea of dimensionless ratios.

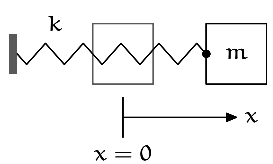

The problem is to find the period of an oscillating spring–mass system given an initial displacement x_0, then allowed to oscillate freely. The relevant variables that determine the period T are mass m, spring constant k, and amplitude x_0. Those three variables completely describe the system, so any true statement about period needs only those variables.

This is pretty neat. I never thought about it like that before.

they should teach this reasoning in high school, dimension thoery is somewhat avoided completely

This is really neat–so then any statement can be written as dimensionless? by just dividing by one of the terms?

Yes, any statement can be rewritten as dimensionless. And yes, just by dividing by one of the terms since all the dimensions match to begin with, dividing everything by the same term will still keep all the dimensions the same for each term and leave it dimensionless.

This is a clever trick and definitely worth remembering.

It's also worthwhile to point out that this works because, by design, everything had the same dimensions from the very beginning. Which could be another potentially useful fact.

–Edit: And reading another two lines down I see that this was brought up. :P

I agree, really cool way to think about problems!

I actually don't see how adding different dimensionless ratios can give any valid information, even if the dimensions work out

I don't understand how adding this ratio to a dimensionless one is proof that it is dimensionless...?

I believe the point is that because the first one has no dimensions, adding the second one and getting a third requires that the final two have no dimensions. It's just a rule of math.

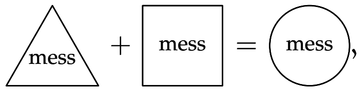

where the various messes mean 'a horrible combination of E, s, ν, ρ, L, and m.

As horrible as that true statement is, it permits the following rewriting: Divide each term by the first one (the triangle). Then

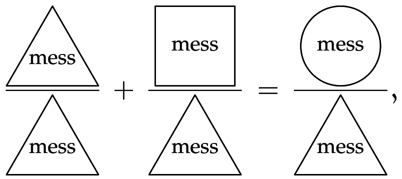

The first ratio is 1, which has no dimensions. Without knowing the individual messes, we don't know the second ratio; but it has no dimensions because it is being added to the first ratio. Similarly, the third ratio, which is on the right side, also has no dimensions.

So the rewritten expression is dimensionless. Nothing in the rewriting depended on the particular form of the true statement, except that each term has the same dimensions.

Therefore, *any true statement can be rewritten in dimensionless form.*

Dimensionless forms are made from dimensionless ratios, so all you need are dimensionless ratios, and you can do all your thinking with them. Here is a familiar example to show how this change simplifies your thinking. This example uses familiar physics so that you can concentrate on the new idea of dimensionless ratios.

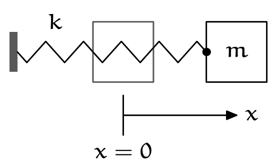

The problem is to find the period of an oscillating spring–mass system given an initial displacement x_0, then allowed to oscillate freely. The relevant variables that determine the period T are mass m, spring constant k, and amplitude x_0. Those three variables completely describe the system, so any true statement about period needs only those variables.

Interesting point. I hadn't realize this until reading this section.

Yep this is really cool.

Hmm, interesting. Can someone give me a quick simple example?

I think the easiest way to think about it is that a lot of true statements that we assert can be expressed as equalities. When you take a true statement, such as F_drag=F_gravity, the dimensions on both sides of the equation are already the same. The left and right side of the equation are in the units of Newtons (dimension would be force). At this point any division of the two sides that preserves the equality also preserves the same dimensions (its like an invariant of equalities). If you simply divided left by right you would have two dimensionless sides that contain the exact same information as the initial equation (force/force = 1).

Why is this true?

Should this be "any true relation among physical quantities"? I don't see how to write the sentence "Mr. X's favorite dimension is the length dimension" in dimensionless form.

And why the restriction to true statements? Surely any false statement can also be rewritten in dimensionless form. Sentences like "the ship's velocity is equal to its mass" are not false, they're meaningless.

> You're right that I should add the restriction "among physical quantities". (Now I had better check whether I said it correctly in _Street-Fighting Mathematics_.)
>
> That's a good point about false versus meaningless statements. I never thought about it that way before, and I think I'll revise the claim to, "any meaningful relation among physical quantities." Although sometimes it's hard to tell the difference between false and meaningless (as the Exxon/Nigeria quote shows).

We've always been taught that units needs to match in an equation. Time can't equal speed. I run into the problem of needing to make things dimensionaless a lot in coding models.

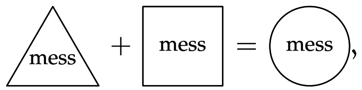

where the various messes mean 'a horrible combination of E, s, ν, ρ, L, and m.

As horrible as that true statement is, it permits the following rewriting: Divide each term by the first one (the triangle). Then

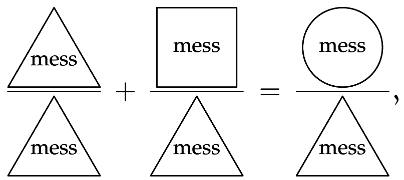

The first ratio is 1, which has no dimensions. Without knowing the individual messes, we don't know the second ratio; but it has no dimensions because it is being added to the first ratio. Similarly, the third ratio, which is on the right side, also has no dimensions.

So the rewritten expression is dimensionless. Nothing in the rewriting depended on the particular form of the true statement, except that each term has the same dimensions.

Therefore, *any true statement can be rewritten in dimensionless form.*

Dimensionless forms are made from dimensionless ratios, so all you need are dimensionless ratios, and you can do all your thinking with them. Here is a familiar example to show how this change simplifies your thinking. This example uses familiar physics so that you can concentrate on the new idea of dimensionless ratios.

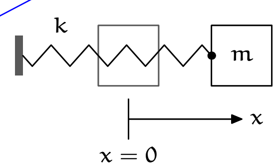

The problem is to find the period of an oscillating spring–mass system given an initial displacement x_0, then allowed to oscillate freely. The relevant variables that determine the period T are mass m, spring constant k, and amplitude x_0. Those three variables completely describe the system, so any true statement about period needs only those variables.

Ok so after the first section, which showed exactly how important the dimensions are, why are we being convinced to forget about them? What's the advantage?

> You won't be forgetting them. You'll use them – to figure out how to write statements in dimensionless form. The dimensions will then tell you the possible structures of that dimensionless statement.

If you are trying to make something dimensionless, is it reasonable to divide the quantity by another quantity with dimensions that you already know the quantity of? Ie. I am trying to make mass dimensionless and I divide it by the mass of a ball that I know.

If you do that, you are never going to get any answers.....if you just divide by the mass of your ball for example, you will just get 1...

> It's good to make quantities dimensionless, but you want to use a relevant mass as the standard of comparison (i.e. as the denominator. So, if you are comparing various sports, it might be useful to list all the ball masses in terms of one of them.

But wouldn't their dimensions effectively become the ball? I.e. a basketball's volume is 2.5 footballs (that was a guess).

I understand from physics why these variables completely describe the system, but how would I know this a priori?

> One way is to think, "How can I describe the system completely, and what is information is required for that description?" Here, the spring differential equation and initial conditions completely describes the system because its solution tells you the entire (past and) future history of the spring.
>
> To write the differential equation, you need to know k and m. To write the initial conditions, you need to know x_0 (the initial velocity is zero by assumption). So k, m, and x_0 are all you need to put into the system. And the period T is what you get out of the system.

I think here you don't have to know it, you're being told. A deeper explanation of _why_ is probably beyond the scope here.

Since any true statement can be written in dimensionless form, the next step is to find all dimensionless forms that can be constructed from T, m, k, and x_0. A table of dimensions is helpful. The only tricky entry is the dimensions of a spring constant. Since the force from the spring is $F = kx$, where x is the displacement, the dimensions of a spring constant are the dimensions of force divided by the dimensions of x. It is convenient to have a notation for the concept of 'the dimensions of'. In that notation,

Var	Dim	What
T	T	period
m	M	mass
k	MT^{-2}	spring constant
x_0	L	amplitude

$$[k] = \frac{[F]}{[x]},$$

where [quantity] means the dimensions of the quantity. Since $[F] = MLT^{-2}$ and $[x] = L$,

$$[k] = MT^{-2},$$

which is the entry in the table.

These quantities combine into many – infinitely many – dimensionless combinations or groups:

$$\frac{kT^2}{m}, \frac{m}{kT^2}, \left(\frac{kT^2}{m}\right)^{25}, \pi\frac{m}{kT^2}, , \ldots.$$

The groups are redundant. You can construct them from only one group. In fancy terms, all the dimensionless groups are formed from one *independent* dimensionless group. What combination to use for that one group is up to you, but you need only one group. I like kT^2/m.

So any true statement about the period can be written just using kT^2/m. That requirement limits the possible statements to

$$\frac{kT^2}{m} = C,$$

where C is a dimensionless constant. This form has two important consequences:

1. The amplitude x_0 does not affect the period. This independence is also known as simple harmonic motion.

There was a long interlude, so now I don't remember what we are doing and what the "next step" is exactly. Maybe clarify just a little in this sentence. The reader shouldn't have to flip back a page to understand your thought process.

These two capital T's are confusing. At first, I had gone through the pages thinking they were the same.

Only later did I realize one was serif and the other sans serif. Perhaps, it would be clearer and more consistent if you used a small "t_p" for the variable time or just something to visually distinguish the two.

agreed

It's hard because the period is usually denoted T and if you're abbreviating dimensions, it seems most straightforward to have them as capital letters. An unfortunate coincidence here, you're right.

This makes me wish you could call my 7th grade algebra teacher and get points back for 7th grade me for not writing units down in my answer.

haha, my 5th grade math teacher would always ask us if the units were aardvarks when we didn't write any down. Although we did just demonstrate the first part of this reading that dimensions are very important!

If your answers were supposed to have units, I don't think Sanjoy would be on your side.

Besides, wouldn't your answers be different numbers if you wrote them in dimensionless ratios?

But I do agree with your point.

Making a little table like this is very useful when performing dimensionless analysis

it is helpful for analysis, but i still dont see the benefit to doing extra work just to find dimensionless quantities.

Well this looks like it could have helped the people building the Mars spacecraft....I think the benefit in working with dimensionless units is that they are more helpful when working with more people.

Since any true statement can be written in dimensionless form, the next step is to find all dimensionless forms that can be constructed from T, m, k, and x_0. A table of dimensions is helpful. The only tricky entry is the dimensions of a spring constant. Since the force from the spring is $F = kx$, where x is the displacement, the dimensions of a spring constant are the dimensions of force divided by the dimensions of x. It is convenient to have a notation for the concept of 'the dimensions of'. In that notation,

Var	Dim	What
T	T	period
m	M	mass
k	MT^{-2}	spring constant
x_0	L	amplitude

$$[k] = \frac{[F]}{[x]},$$

where [quantity] means the dimensions of the quantity. Since $[F] = MLT^{-2}$ and $[x] = L$,

$$[k] = MT^{-2},$$

which is the entry in the table.

These quantities combine into many – infinitely many – dimensionless combinations or groups:

$$\frac{kT^2}{m}, \frac{m}{kT^2}, \left(\frac{kT^2}{m}\right)^{25}, \pi\frac{m}{kT^2}, , \ldots .$$

The groups are redundant. You can construct them from only one group. In fancy terms, all the dimensionless groups are formed from one *independent* dimensionless group. What combination to use for that one group is up to you, but you need only one group. I like kT^2/m.

So any true statement about the period can be written just using kT^2/m. That requirement limits the possible statements to

$$\frac{kT^2}{m} = C,$$

where C is a dimensionless constant. This form has two important consequences:

1. The amplitude x_0 does not affect the period. This independence is also known as simple harmonic motion.

unless F has a unique coordinate system, or it's made clear that we are only discussing magnitudes, I feel like the minus sign shouldn't be omitted from F=-kx (a system with F=kx would have no periodic oscillation given k>0)

I don't think the sign matters because we are looking at the dimensions of the formula not the final answer.

May be good to explicitly say that the notation is brackets around the symbol for the quantity.

I really like this way of determining the units.

This is only a bit confused because of the odd use of brackets earlier!

I think they were in a different context

Var	Dim	What
T	T	period
m	M	mass
k	MT^{-2}	spring constant
x_0	L	amplitude

Since any true statement can be written in dimensionless form, the next step is to find all dimensionless forms that can be constructed from T, m, k, and x_0. A table of dimensions is helpful. The only tricky entry is the dimensions of a spring constant. Since the force from the spring is $F = kx$, where x is the displacement, the dimensions of a spring constant are the dimensions of force divided by the dimensions of x. It is convenient to have a notation for the concept of 'the dimensions of'. In that notation,

$$[k] = \frac{[F]}{[x]},$$

where [quantity] means the dimensions of the quantity. Since $[F] = MLT^{-2}$ and $[x] = L$,

$$[k] = MT^{-2},$$

which is the entry in the table.

These quantities combine into many – infinitely many – dimensionless combinations or groups:

$$\frac{kT^2}{m}, \frac{m}{kT^2}, \left(\frac{kT^2}{m}\right)^{25}, \pi\frac{m}{kT^2}, , \ldots .$$

The groups are redundant. You can construct them from only one group. In fancy terms, all the dimensionless groups are formed from one *independent* dimensionless group. What combination to use for that one group is up to you, but you need only one group. I like kT^2/m.

So any true statement about the period can be written just using kT^2/m. That requirement limits the possible statements to

$$\frac{kT^2}{m} = C,$$

where C is a dimensionless constant. This form has two important consequences:

1. The amplitude x_0 does not affect the period. This independence is also known as simple harmonic motion.

I did not know this relation, where did it come from?

F = ma where m has dimensions of M and a has dimensions of L/(T^2). (Think about common units for acceleration: m/(s^2).)

maybe you could explicitly define M, L, and T as mass, length, and time somewhere? Especially to distinguish M from m... (Mass in general from a specific mass from meters... I think we need more letters

This got pretty confusing too...why do so many things start with m?

He does define them in that box at the top of the page.

That's true. But in terms of clarity, I am a proponent of the easier it is to decipher, the better it is in terms of benefiting the reader.

In this case, perhaps being more literal by writing out [F]=[Mass][Length][Time]^-2 would help since that is the most basic starting premise.

Then, moving to substitute [time]=[Period] with the justification that the magnitudes are meaningless.

Lastly, specifying the relations for [x]=?, [k]=? and how they match the parts of [F].

Then say that you are choosing the [k]= relation to form a dimensionless combination.

Another thing to consider is whether if it might be more appropriate to use "proportional" signs instead of equal signs in [F]=, [x]=, and [k]=. This way, the C is not so unexpected, when it shows up.

This comment and several others have got me thinking about adopting the following policy in the book: No abbreviations! Edwin Taylor, who wrote the best special-relativity textbook (_Spacetime Physics_) ever written, told me that he is doing that for the revised edition of his general-relativity textbook (_Exploring Black Holes_).

The idea intrigued me as soon as I heard it, and the more comments I read, the more convinced I am that it would be a good idea. In short, why create roadblocks to understanding, even small ones?

Since any true statement can be written in dimensionless form, the next step is to find all dimensionless forms that can be constructed from T, m, k, and x_0. A table of dimensions is helpful. The only tricky entry is the dimensions of a spring constant. Since the force from the spring is $F = kx$, where x is the displacement, the dimensions of a spring constant are the dimensions of force divided by the dimensions of x. It is convenient to have a notation for the concept of 'the dimensions of'. In that notation,

Var	Dim	What
T	T	period
m	M	mass
k	MT^{-2}	spring constant
x_0	L	amplitude

$$[k] = \frac{[F]}{[x]},$$

where [quantity] means the dimensions of the quantity. Since $[F] = MLT^{-2}$ and $[x] = L$,

$$[k] = MT^{-2},$$

which is the entry in the table.

These quantities combine into many – infinitely many – dimensionless combinations or groups:

$$\frac{kT^2}{m}, \frac{m}{kT^2}, \left(\frac{kT^2}{m}\right)^{25}, \pi\frac{m}{kT^2}, \ldots.$$

The groups are redundant. You can construct them from only one group. In fancy terms, all the dimensionless groups are formed from one *independent* dimensionless group. What combination to use for that one group is up to you, but you need only one group. I like kT^2/m.

So any true statement about the period can be written just using kT^2/m. That requirement limits the possible statements to

$$\frac{kT^2}{m} = C,$$

where C is a dimensionless constant. This form has two important consequences:

1. The amplitude x_0 does not affect the period. This independence is also known as simple harmonic motion.

very well explained, thank you.

though i agree that this is most certainly correct, i find that, in practical problem solving, it is usually far more helpful to say things like [x]=kg, and [k]=kg*s^-2. this helps prevent problems such as that of the mco.

> I also find myself referring more often to the units being used rather than general dimensions like Mass or Time. Are there any disadvantages to taking this approach instead?
>
> > Scale could be a problem, if you're thinking in milligrams about an elephant. As in the previous unit, it's generally better to think "if this system had more mass, then x would happen" instead of "if x had more kilograms." Force is mass times distance over time squared, not a kilo times a meter over squared seconds (that's a Newton, not force, and if you picked other units, you might be thinking about some unusual unit no one uses) Try just using the particulars of the units to remind you which dimensions you need (oh, a Newton has kg in it, so that's mass...) and then cast aside the particular units.

I really like how you were able to make dimensional analysis systematic. However, sometimes it's hard to pick the powers just by inspection - sometimes you can make combinations involving things to the 2/3 power and what not. It'd be nice to learn a systematic way to do that part too

You might want to be more specific...'The quantities reliant to this problem are, however, redundant: you can construct them all from just one group."

So then when you find these, what do you do next? you just rewrite the equation in terms of this "dimensionless" unit?

I'm confused by how T went from something that seemed like a dependent variable of the 3 things that completely defined the system (above) to now being on the same footing as those variables when we look for dimensionless combinations. I knew beforehand that xo didn't affect the period, but why is THIS the outcome of our finding that it's not in the dimensionless constant? (devil's advocate: why not say that xo doesn't affect the mass?) You sort of cheat to get there by not including it in the list, since kxT^2/mx is also a valid ratio. There could be a comment about combining things in the simplest way possible.

Since any true statement can be written in dimensionless form, the next step is to find all dimensionless forms that can be constructed from T, m, k, and x_0. A table of dimensions is helpful. The only tricky entry is the dimensions of a spring constant. Since the force from the spring is $F = kx$, where x is the displacement, the dimensions of a spring constant are the dimensions of force divided by the dimensions of x. It is convenient to have a notation for the concept of 'the dimensions of'. In that notation,

Var	Dim	What
T	T	period
m	M	mass
k	MT^{-2}	spring constant
x_0	L	amplitude

$$[k] = \frac{[F]}{[x]},$$

where [quantity] means the dimensions of the quantity. Since $[F] = MLT^{-2}$ and $[x] = L$,

$$[k] = MT^{-2},$$

which is the entry in the table.

These quantities combine into many – infinitely many – dimensionless combinations or groups:

$$\frac{kT^2}{m}, \frac{m}{kT^2}, \left(\frac{kT^2}{m}\right)^{25}, \pi\frac{m}{kT^2}, , \ldots.$$

The groups are redundant. You can construct them from only one group. In fancy terms, all the dimensionless groups are formed from one *independent* dimensionless group. What combination to use for that one group is up to you, but you need only one group. I like kT^2/m.

So any true statement about the period can be written just using kT^2/m. That requirement limits the possible statements to

$$\frac{kT^2}{m} = C,$$

where C is a dimensionless constant. This form has two important consequences:

1. The amplitude x_0 does not affect the period. This independence is also known as simple harmonic motion.

in 2.006 I thought there was a slightly more formulaic way to find what the dimensionless groups were- any suggestions?

I like this way of thinking - looking forward to this unit.

I agree, this is some pretty eye opening stuff.

So theses dimensionless constants also help us see more simplicity in the problem?

Is making something dimensionless also kind of like abstracting?

Yes! The dimensionless quantities are themselves abstractions. For example, the famous Reynolds number, a combination of speed, size, and (kinematic) viscosity, is dimensionless – and it hides all the details that we don't care (like the particular speed or size or viscosity). All we care about (in some situations) is the Reynolds number itself.

I think this second sentence should read 'This independence is a property of simple harmonic motion' for two reasons:

- Simple harmonic oscillation is usually defined by the system obeying Hookes' law (F=-kx) (and not usually by the result that the amplitude is constant)

- Technically, simple harmonic motion is not the only way to have this independence (i.e., the independence of amplitude and period can exist for other types of systems).

I think you should point out the usefulness of this tool in it's ability to determine that it is simple harmonic motion. Sure, it's nice to use dimensionless analysis, but it is also good to be able to omit dimensions from a problem to understand what is causing changes in the system.

2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here $h()$ means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [19, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

Do you mean it is independent of both of them, or that it only depends on their ratio and not on them individually?

It only depends on the ratio and not each of them individually. Sort of like a lot of the examples we have been doing this past unit.

Why not also say independent of T? The useful outcome of this that I'm most used to is the formula T sqrt(m/k), and it might be nice to bring this up as a result.

2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here $h()$ means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [19, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

I'm missing something.

Right now, I am thinking that C is constant for a specific set of k, m, T. Why would a different set of k, m, T necessarily have the same value of C?

How does that follow from above?

Ok, I think I get it now.

You started from the equality:

[F]=MLT^-2

and matched the different dimensions with their variables.

Then rearranging the variables to one side, you'll be left with a unitless unknown factor, C, that accounts for not worrying about exact relations when focusing only on the dimensions.

Since the force relation is an equality, every set of k, m, T must obey that and thus, C will be constant.

There are 2 things that impeded my understanding initially:

1) no visual representation of the mapping/substitution when matching the dimensions of k and m into the dimensions of F.

2) using period instead of time within the dimension of [F] itself. I thought it was a nonspecific time at first and also the two similar looking letters for the variable and the unit helped hide it all the more. This is rather a sneaky/subtle maneuver, because on first thought, I wonder whether it is appropriate to use Period in the units for Force instead of a generic "time, t". But then, I realize it won't matter, since the magnitude of the period, T, will be absorbed/offset by the extremely accommodating unknown constant, C.

This is really cool and challenged my original preconceived ideas, especially the inflexible way I was taught of adhering to classical representations of these relations. The dimensional analysis approach eliminates the inflexibility of exact relations by extracting useful ideas in the form of a group of familiar dimensions and shifting all the trivial exact magnitudes onto a constant, C.

shiny! very useful to know...thanks

Why is C independent of k and m? Aren't they in the equation for it?

2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here h() means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [19, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

This whole statement is confusing. Aren't k and m in the constant definition?

So using this type of dimensional analysis actually gives insight to problems–I never would've thought so! Before I always thought dimensional analysis was just for checking your answers.

Indeed it is simpler. I wish I would have had this insight when I was solving physics problems earlier in my academic career.

> I guess it might be easier to figure out the dimensions of a problem especially physics before you tangle the problem...It would definitely make it easier to think more clearly about the equations I might want to use.

I'm not really sure that defining h as a function of no variables instead of just stating that it should be a constant really adds anything to the discussion.

I think you may have mixed up left and right, or else I am very confused. If you were referring to the right side of the equation T = f(m, k x0), then it might make more sense. This whole section could be much clearer.

are you saying that h is a function of 0 variables because all the variables are already being used on the Left side? so if you didn't use one of the variables would the right side be a function of that? and both sides are always dimensionless correct?

variables is not a concrete term, a variable is something that you don't know, just because it has no variables in this case doesn't mean it always is variableless

Right side of which equation?

> It is confusing as it is worded. It means the equation presented immediately above, the right side of which is h(). If the right side were to have variables, they would have dimensions k, T, and m (no other dimensions are relevant here). However, we know by dimensional analysis that they would have to take the form kT^2/m, which is already on the left side. Thus it can be said that h() depends on nothing.
>
> The fact that he says "h() depends on no variables", and then says: "the right side contains all other quantities" seems contradictory, as worded.
>
> I'm not sure about this either.

2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here h() means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [19, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

"Why no variables?" would make a good side bar ... you only need to read it if you don't understand why h() is a function of no variables. [i read it thinking...yeah, yeah, come on, next]

I'm still a little confused. So I understand the example, but I don't see what situations dimensionless groups would be useful in. Since dimensionless groups can represent (what is seems like) many different things, how does it help us in finding answers to our calculations?

An easy useful dimensionless number is the Reynolds number, and depending on the number, you know if the flow is turbulent or laminar. Also, to find the heat transfer coefficent as a function of known properties, you use the Nusult number. In this case, the point is if you know C, and you know T and k, you can solve for m no problem.

I understand how you made things dimensionless but not what the solution of the problem. Does this mean that the period is a constant?

It might be worth it to mention that C turns out to be 4*pi^2 and say as an exercise prove this through the physics.

what era was he most prominent? using measurements such as bookcases and libraries seems to date him.

> Roughly, the whole of the 1900s. He was a brilliant physicist and statistician. Alas, his approach to statistics, the Bayesian approach (pioneered by Laplace), got overshadowed by the biologists' approach (p-values, confidence intervals), which is now called "orthodox statistics."
>
> But the winds are turning toward Bayesian statistics, and I'll try to show the fundamental ideas in the unit on "probabilistic reasoning" (part of the next group of methods on how to discard information).

To prior to the last 10 years? People still have all these items in physical form, and we still use these concepts as metaphors all over the place (web pages, DNA libraries...).

what is meant a table of functions of one variable? is it a table of every function that that variable is in?

2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here $h()$ means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [19, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

I ran across this quote when I was in high school, and estimated that he wasn't too far off!

That's a cool quote, it really helps add to the argument here about the power of dimensionless quantities. I'm glad its in the text.

How would this work with equations whose forms depend on dimensions? For example, Maxwell's Eq look different for CGS units and SI units.

I think they look different for different definitions of units, not that the forms depend on dimensions per se.

Er, wait, I'm confused. Are we trying to find dimensionless invariants, or are we trying to use dimensions...?

That phrasing needs improvement. I meant, use dimensions to find dimensionless quantities. And these quantities are invariants in that they are independent of the system of units (e.g. if you change from meters to furlongs, you don't change the Reynolds number).

This explanation makes a lot of sense. The next paragraph also ends up clearing a lot of the question up.

sounds like a concept in group thoery

This is a really nice paragraph tying together a lot of what we've learned and really makes me think about how I should be approaching problems.

I think what becomes key when using dimensionless analysis is that we solve for an INVARIANT.

This definitely clarifies matters in terms of how this ties in with the previous sections.

I agree, this paragraph really ties past units together nicely! before I was confused as to how we were supposed to know what invariants to look for, now I realize that looking at units is a really great way to define them. However, some invariants that depend on numbers alone (the cube game, for example) would still be difficult to find...

2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here $h()$ means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [19, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

Are we going to develop a method for looking for dimensionless groups? We did in 2.006 and I'm curious how this class will do it.

> We will, but you won't be surprised to know that I hardly ever use or teach the formal, linear-algebra method (where you solve simultaneous equations). Instead we'll find them by educated guessing.

"is a recipe"

Before reading this section, I never really thought much about dimensional analysis beyond being something that I could use to check my answers. Now, I'll try to use them to find answers or simplify problems before I kill myself trying to solve something harder than it should be.

Yeah! My thoughts exactly. This makes dimensional analysis much more useful than I initially thought.

I completely agree! It definitely help to see how all the items of this class are important and in some how related. Initially when looking at the different forms of approximations we have used I assumed after this class I would only use divide and conquer. However I am starting to see how everything is related and see the significance of most of the approximation techniques we have used so far.

This somewhat occurred to me back when we did abstraction, and I'm glad to see my hunch was correct

Any quick examples?

this is very much in line with proportional reasoning, "look for what does not change"

yeah this is a cool way to build on that approach

Yeah it really shows why we did proportional reasoning before this section. I like it.

2. The constant C is independent of k and m. So I can measure it for one spring–mass system and know it for all spring–mass systems, no matter the mass or spring constant. The constant is a universal constant.

The requirement that dimensions be valid has simplified the analysis of the spring–mass system. Without using dimensions, the problem would be to find (or measure) the three-variable function f that connects m, k, and x_0 to the period:

$$T = f(m, k, x_0).$$

Whereas using dimensions reveals that the problem is simpler: to find the function h such that

$$\frac{kT^2}{m} = h().$$

Here $h()$ means a function of no variables. Why no variables? Because the right side contains all the other quantities on which kT^2/m could depend. However, dimensional analysis says that the variables appear only through the combination kT^2/m, which is already on the left side. So no variables remain to be put on the right side; hence h is a function of zero variables. The only function of zero variables is a constant, so $kT^2/m = C$.

This pattern illustrates a famous quote from the statistician and physicist Harold Jeffreys [19, p. 82]:

> A good table of functions of one variable may require a page; that of a function of two variables a volume; that of a function of three variables a bookcase; and that of a function of four variables a library.

Use dimensions; avoid tables as big as a library!

Dimensionless groups are a kind of invariant: They are unchanged even when the system of units is changed. Like any invariant, a dimensionless group is an abstraction (Chapter 2). So, looking for dimensionless groups is recipe for developing new abstractions.

Could we get an example of how we would use this to actually solve a problem?

I'm guessing those will be forthcoming, but the process as I remember from 2.006 is to get your dimensionless groups and then go to the lab and run experiments to see how they are related to each other. Once you have the relationship you can extrapolate to find a lot of other stuff.

"All good things come to those who wait." The upcoming sections, plus the lecture examples, will I hope answer your question.

I particularly enjoyed this chapter and am looking forward to the rest of the units. I used dimensional analysis a lot in a high school physics competition when time was running out and I just starting guessing answers!

And it's also saved me from writing a wrong answer when I realize my units don't match...

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world. . .

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

Another tiny bug you might want to consider taking out of NB is, when you misspell a word, the little red line comes up under the word to tell you. This is nice, except usually when this happens you can right click it and it will show you the correct spelling. It's not a huge glitch, but it'd be nice I think.

I don't see any red lines under misspelled words, so I think it's a browser/platform-specific issue. At least it's telling you when things are wrong.

This definitely took a couple readings to understand completely. I feel like it is a concept fundamental to further exercises, so I won't suggest that it should be softened, but it is definitely one of the more difficult logical processes.

I feel this way too, it was definitely one of the hardest, if not the hardest, sections to read through. Given the depth of the calculations, I feel that it lacked helpful summarization and conclusions of the concept which the previous chapters had done well one.

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world...

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

Read this section for Thursday's memo. Don't worry about the funny page breaks. (All will change in the summer when I sit down with all your comments...)

so

Might it be too simple to understand all the essentials of atomic theory though?

Yes, but for a lot of things it is realistically the ONLY one we can study. The rest, even helium, are just too complicated

As far as I remember, it's the only atom we studied in quantum mechanics I. And even that was complicated...

Also, one of our mantras is to start simple, try to get somewhere, and then add complexity.

oooh, I like Feynman lectures, great use of ethos.

If all knowledge were to be destroyed, I really doubt that the future people would believe this statement at face value...

interesting to think about

I wonder what a famous scientist of another background might pick.

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world. . .

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

This definitely isn't the one thing I would've thought of, but I also don't know the enormous amount of information that can be found in that one sentence.

It definitely is a core concept to our understanding of the world and universe that we probably take for granted though. It is hard to imagine not having this piece of information in our minds.

It's what all the laws of of physics reign over.

Wouldn't instructions be more helpful?

This seems so informal. But I guess with the going of scientific knowledge, so goes the scientific jargon.

I guess i agree with this statement–everything is based on atoms after all right...it's what everything is made of and the interactions between atoms go a long way in explaining almost everything in the physical world...

Too bad humans have this knack for not believing things until they are proven.

Will we find out why?

Is it bad that my first thought was to relate this to information theory, and to relate the number of bits required to transmit this information to the total amount of information stored?

Nah, it's natural.

Did he just postulate this or did he have some proof?

Yeah, how did he end up hypothesizing this with no way of measuring or seeing them? That takes some serious insight..

Oh, I think they had all sorts of educated guesses back then. Perhaps he just happened to be the lucky one who wrote his thoughts down.

'Who is Democritus?" would be a good side bar note...I'm apparently not letting go of this idea. sorry, if it bothers you

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world...

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

Has anyone read this? Is it worth reading?

I've read it, and I would say it's worth the read if you have the time. It gives a great overview of the philosophers and theories that came before us.

I've read parts of it. It's interesting but long, if I remember.

Bertrand Russell has a lot of good works.

As far as I know, Bertrand Russell was a great philosopher but not a great historian of philosophy. The book is fun to read but don't trust it to get the views of past philosophers right.

I think this is like the Reader's Digest version of Western Philosophy.

However, with...

Coming after "classical and quantum mechanics", dimensional analysis seems...quaint. All we know about atoms is described by quantum and classical mechanics, so it seems like dimensional analysis is just tacked on because that's what you want to talk about.

I agree. I think if you know classical and quantum mechanics then dimensional analysis probably isn't particularly necessary.

hm maybe dimensional analysis can serve as an easier way to think of classical and quantum mechanics?

This transition seems really forced. From the intro it sounds like you are going to explain atomic theory using mechanics (because that is the basis behind the really broad topics that are talked about in the quoted passage) but in the next section the focus is on dimensional analysis.

well the whole chapter is on dimensional analysis, so we knew it was coming

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world...

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

Do you think we could use this for every problem on the "diagnostic pretest"? Why do we need other methods when this one does so well?

I'm not sure what you mean by using Dimensional Analysis for every problem. Is that possible? I am still understanding the Dimensional Analysis. Care to explain?

It's always useful, but more so to check your answer or give you an equation. I don't see the help in estimation, like the sequences or proportionalities.

> The sequences on the diagnostic is a good example of where dimensional analysis cannot help you. The goal is a pure number (the size of the nth term) as a function of a pure number (n). Both items are dimensionless, so any function would be dimensionally okay. That is, dimensional analysis doesn't place any restriction on what the function is. You need other techniques to make a full toolbox.

This whole concept seems very useful, if mastered. It could significantly help on tests to use units (I don't but probably should)

I don't know how useful this is a heading though... isn't the whole section on that.

This class integrates so many different aspects of Science and I think that is so interesting.

Do you mean to say "uses the hydrogen atom..."?

Agreed, it sounds awkward as written.

This first sentence is also a bit awkward, mostly because of the "in order...two questions" phrase at the end.

is "it" referring to an atom?

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world...

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

For some reason I found this phrasing a little confusing and had to do a double-take...

I agree. The phrasing is confusing and breaks up the flow of the sentence.

This and the previous sentence could be combined to keep the flow going - at the moment it brings the flowing style of the prior readings to a staccato pattern.

i think the confusion has to do with the ambiguous "it" and "that"

Perhaps, something along the lines of the following:

"The first question is how big it is, which sets the standard for the size of more complex atoms and molecules"

Why does it set the size for more complex atoms?

Because the more complex atoms are built from the smaller atoms.

I would've liked to see some bullet points. Clears things up. Especially when your asking a 2 part question.

Why are these questions important? or rather why are they so important?

I know it states that it let's you find size of more complex, and things like stiffness, speed and energy content. But without understanding some fairly complex physics that relates bond energy to these things, how do we go about relating things like that on a daily basis?

are we going to take the periodic table into account?

Should this better read "how big it is".

The "it" could be a bit ambiguous. Are we trying to find radius, volume, something else? (I realize they are all related)

I think he's referring to how big is the atom itself (dimensionwise)... but this sentence definitely reads awkwardly.

"how big it is" wouldn't be a question.

I am assuming this will be in a future section? I did not see it in this one.

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world...

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

Disassemble it nuclearly - or its molecular bonds? Are we just looking at the scope of the Atom

hydrogen is an atom, not a molecule, so there are no molecular bonds. To clarify, perhaps the clause "into its subatomic particles" could be added?

Since there aren't any bonds, how do you disassemble it? and for what reason? I think the termiology is confusing me.

I guess we could disassemble to protons and electrons.

I think just saying "separate the electron from the proton" would make things a lot clearer.

Hydrogen is diatomic as a gas, so it could be the separation of one hydrogen atom from the other in H2. That would be a bond energy as mentioned in the next section. It could also be nucleus/electron separation, which is what happens when it's ionized.

Couldn't it also mean separating the proton from the neutron? I'm not sure why you would want to do that but....

Each lecture always starts off so randomly- but they all link together! Again, you're using the most basic element (literally) to calculate more complex ones

this whole paragraph is awkwardly written.

This wasn't immediately obvious to me, but perhaps I've forgotten chemistry.

I think just means that since hydrogen atoms are prevalent in many complex substances, it's useful to know the energy of a hydrogen bond.

Agreed. I don't think he's saying is the absolute basis of all bonds, just that it's a good basis for comparison and is very important in chemistry and physics

i will argue this point... there are a lot of factors that determine stiffness. this would only be true if we were talking about an absolute pure material in crystalline form

I agree, since all materials are made up of hydrogen-like atoms, they should all have relatively similar bond strength. However, this is not the case even with the same elements. Graphite and diamond are both made from carbon and have very different bond strengths. I'm not sure how plausible this conclusion is without mentioning that there are other factors involved.

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world. . .

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

Nice foreshadowing here... gets the reader intrigued about the powerful capabilities of dimension analysis that will soon be explained.

I agree, it makes an impact showing a group of things determined by these energies that are not all completely related.

I agree, it makes an impact showing a group of things determined by these energies that are not all completely related.

I like knowing that there is a reason for learning what I'm learning... that it is useful somehow and applicable to the real world.

I really like this sentence as a way to get the reader excited about why what he or she is about to read is very useful and can be used to derive many important quantities.

I'd use: "All from the basic understanding of hydrogen!"

well it's probably somewhere around half.

what are you referring to?

My earlier statement, about stiffness not solely being dependent on the scaling of bond energies from hydrogen, makes me question the other conclusions like speed of sound and energy from sugar. Its just hard for me to believe that things are that simple.

I would assume that several other factors come in to play to determine these quantities, although we can probably still learn a lot from hydrogen. To me it seems a little too ambitious...

I agree with both those statements, but I think that going straight to the equation takes away from learning how to get variables. Maybe add a bit more and start with a diagram (we have radius, chargy, energy, etc). When it went straight to the equation, I was thinking "I would have never have gotten that right off the bat."

add: "approach"

Is this the first step in answering the question of how big it is?

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world...

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

Would be useful to have summaries at the end of each chapter w/ a worked through example, at least in the final version.

I thought we were working through examples in the text? Would you like more problems to practice?

What would a more complex model of hydrogen be? I've only ever seen it modeled this way.

You stated that we wanted to find the energy to break an atom, then gave us the equation for the force, then from that got the necessary parameters to get the energy. It seems all very backwards and contrived. Maybe if you figured out another way to get the necessary parameters, and from that got the force or energy, it would make the whole example a lot more legitimate.

Here you use e for the elementary charge, and later you talk about a charge q without relating the two. I feel like this sort of obscures your point and you might want to make a note of this, especially because a casual reader might not recognize that e represents the charge of an electron.

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world. . .

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

I thought 8.02 was not a prereq for this class? You might want to mention kqq/r^2 to jog people's memory.

I think this equation is simple enough that a relatively advanced high school student could grasp it (and I think high school and above is the target audience).

I agree with it being simple - while the symbols might initially be scary you can break it down to some constant and r^-2

It's not a pre-req, which is why the equation is given here. It's pretty intuitive because all the variables are explained.

8.02 would be a pre-req if instead hte book just said "the magntiude follows the electrostatic force equation" or coulombs law.

I'm pretty sure most of these equations were taught in high school physics class.

You'd be surprised.

I didn't take physics in high school so I hadn't seen it before 8.022. But I agree, it's pretty straight forward and should be left as is.

and we're more focused on learning the method of dimensionless analysis then derivations.

Adding F = this would make it more standard-looking.

Certainly nothing would be lost by adding in the general form for the electrostatic force, and it would probably help make this section more clear, even to people who remember 8.02 well.

it's essentially the same as kqq/r^2. I think if you've seen it before and didn't remember it, you could just look closely at the equation and realize that, oh, it's just kqq/r^2.

5.3 Hydrogen atom

Hydrogen is the simplest atom, and studying hydrogen is the simplest way to understand the atomic theory. Feynman has explained the importance of the atomic theory in his famous lectures on physics [9, Volume 1, p. 1-2]:

> If, in some cataclysm, all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations of creatures, what statement would contain the most information in the fewest words? I believe it is the *atomic hypothesis* (or the atomic *fact*, or whatever you wish to call it) that *all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.* In that one sentence, you will see, there is an *enormous* amount of information about the world. . .

The atomic theory was first stated by Democritus. (Early Greek science and philosophy is discussed with wit, sympathy, and insight in Bertrand Russell's *History of Western Philosophy* [26].) Democritus could not say much about the properties of atoms. With modern knowledge of classical and quantum mechanics, and dimensional analysis, you can say more.

5.3.1 Dimensional analysis

The next example of dimensional reasoning is the hydrogen atom in order to answer two questions. The first question is how big is it. That size sets the size of more complex atoms and molecules. The second question is how much energy is needed to disassemble hydrogen. That energy sets the scale for the bond energies of more complex substances, and those energies determine macroscopic quantities like the stiffness of materials, the speed of sound, and the energy content of fat and sugar. All from hydrogen!

The first step in a dimensional analysis is to choose the relevant variables. A simple model of hydrogen is an electron orbiting a proton. The orbital force is provided by electrostatic attraction between the proton and electron. The magnitude of the force is

$$\frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2},$$

This would look a lot neater in cgs...

cgs?

centimeters, grams, seconds

While I am a big proponent of cgs for E&M problems, I have to say that SI gets the point across better than cgs for people who are not physics majors because everyone is familiar with SI.

Very true. The only E&M experience I have ever had has been with cgs and seeing SI things is strange. It's probably best to pick the unit system that most people are familiar with, regardless of how phenomenally stupid that unit system is.

This kinda comes outta nowhere and there's little detail about it...maybe for non-physics students it might be nice to just throw a sentence or two explaining it's meaning or something

I feel like this is a pretty basic formula and that the majority of the class, even non-physics students like myself, have seen this whether it be in high school or 8.02. Besides, the variables are explained in the few sentences after.

Agreed- right now this just looks like a bunch of variables with no meaning to me. It'd be nice to see a derivation or some background of where this is coming from.

Would it help if it read F=(what is there?

I think it would help if it said "F=" and if it had the name somewhere. "Coulomb's law" or something just so people can realize where they've seen this before.

Are we expected to follow the math completely here?

I think it should be mentioned here, in the last reading, and maybe in the beginning of most examples to not just choose relevant variables, but to also make a table. The table helps immensely.

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = [r^2] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

Comments on page 2

are you going to explain the rest of the variables?

which list? the variables from the thing right above on p1 or this boxed list below? i'm guessing it's the list below but just the word "below" would help

Or the "list of variables we need to think about"

Thought this phrasing was a little awkward but I understand what you mean.

I had to re-read it a few times to understand this phrase.

Using the word 'variable' twice makes it a bit strange. Maybe replace the first one with 'the list below'?

Yeah...I've re read it a few times now and I'm still having trouble understanding it...

You could eliminate e0 with a clever choice of conventions (a la 8.022).

This paragraph doesn't seem entirely clear. I'm not sure how we make the arguments about when q and epsilon naught appear.

are these defined? I don't see them defined.

Can you at least say, explicitly, what these are? I assume q is charge and r radius, but I don't know or remember about epsilon.0. Or what 'e' is, for that matter.

Epsilon_0 is a constant of permittivity, which is the measure of how much resistance is encountered when forming an electric field in a vacuum- its value is about 9*10^-12 F/m

Is q=e?

I didn't follow exactly why since q is relevant that it produces a force, that the charge appears only in the combined equation given.

Nor did I. I also don't understand exactly why it is we should include enough variables to generate the force expression. Is it just because we have to start *somewhere*?

I understood this logic in the case of g*sin(Theta) in class, but not here.

Is this quantity produced dimensionless?

Nevermind, its dimensions are determined later on.

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = \left[r^2\right] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

It would be nice to see a definition of each before getting into this equation. It would at least be nice to see a description of the equation.

I might be biased since I'm an EE but I think we're getting too hung up on where equations is coming from and not on the point of using dimensional analysis. Someone mentioned if 8.02 should be a prereq for this class, but I mean really 8.02 should've been taken freshmen year.

It would be nice to see a definition of each before getting into this equation. It would at least be nice to see a description of the equation.

It would be nice to see a definition of each before getting into this equation. It would at least be nice to see a description of the equation.

I agree, and what is q representing?

I agree, I'm a bit confused with the relevance of each variable.

The table is on the side rather than very integrated in the text. I think if most people could just ignore it and keep reading it. I don't think it's that big of a deal to have it in and someone will find it useful.

Idon't quite understand what this is trying to say

I agree..I'm not quite sure what we're after given the above equation. It seems like everything is there already.

I don't understand what these sentences are trying to explain.

It would be really nice if you mentioned what the units of epsilon_0 are just in case the reader, like myself, has never done E&M in SI units.

It would be really nice if you mentioned what the units of epsilon_0 are just in case the reader, like myself, has never done E&M in SI units.

while e and e_0 may be obvious to a physicist or a chemist,I haven't seen these variables for years. Please define them.

thanks for putting this in, I had forgotten what a Bohr radius was

What makes it ideal?

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = [r^2] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

After reading this whole page, I remembered that in the last reading, you made this list of variables and used it to construct a dimensionless value. However, I didn't remember that at first, so initially I was confused about why we were arbitrarily making this list. Perhaps you could put in a reminder sentence or make a bigger point in the first part from the last reading that making this list is a required step?

I really like the tables... keep them in. It helps me to see quickly what is important and breaks up the text

I would agree with the first comment in this thread.

It is very tempting to read in a linear, up-down fashion. Having the box on the side, decreases its importance in the context of a smooth read. That is, I imagine it is easier to keep reading, rather than to stop the text and carefully examine that table on the side.

Maybe placing that table on its own line, between paragraphs would cue readers to actually look at it the first time through. That way, the table is "connected" within the text, instead of being some reference thing to the side.

yes or maybe adding a caption under the table would be useful.

I still don't understand why this table was included. It actually kept distracting me from the other variables that you were talking about in the reading. I feel that the table on page 93 would have been more suitable here.

Out of every section so far I think dimensional analysis has been my favorite. It seems to be the most quickly applicable and yields the most interesting results, in a way that we would normally never go about solving a problem

I feel like this is the wrong chart for this section...shouldn't the chart here include e & epsilon & r?

After reading this, I understand what all of the variables mean, but I never would have been able to come up with them on my own (at least not all of them). I think this would put a big damper on the accuracy of my analysis - how do I avoid leaving things out?

So why doesn't it?

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = [r^2] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

Why does this table include things like surface tension (it doesn't seem to appear in our analysis)? Our is this just a reference for common dimensions used in general?

A yank (ctrl-Y in Emacs) bug. I copied it from a later section (on waves), for the template, but didn't update it to use the variables for this problem. Whoops.

Whew. I was looking for this comment; I was very lost trying to figure out where all these figured into atoms.

so simple! I actually get it.

I really like this process, it seems so elegant.

i'm not so sure about it actually. it seems to me we're doing a lot of tricks in order to "find" something we should have known the instant we wrote down the equation.

These are interesting units for this equation. The time unit is usually s^-2, but I guess frequency and 1/s are the same.

T means units of time which are seconds here.

I might be mistaken, but didn't you say in the previous reading/section that you can always make things dimensionless with clever manipulations?

Yeah, I agree, I definitely remember that from the previous reading. And I'm still confused... why can't we make a dimensionless value here?

Keep reading! there are no other time values to cancel

You don't have to use all the variables if something won't cancel out.

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = [r^2] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

I like this explanation - it makes the point very obvious.

Which 3 items are you reffering to? q, r and e_nod? How is not being able to make the force dimensionless important?

I mentioned that the equations didn't need to be explained, but I think the dimensions of the different variables should be given. Especially since the topic is in dimensional analysis. The seconds actually comes from epsilon, which can take the form of several units (farads/meter, J/v^2, Amps*seconds/V, etc.) In this case I'm guessing the units for epsilon is given by seconds*Coulombs^2/meters^3*kg. then the coulombs cancel out, leaving out the dimensions given.

I also like the explanation here. I wouldn't have seen that they were not going to make a dimensionless group at first though; I think a proper table would have helped with variables. It is MUCH clearer when there is a table of variables.

isn't the easiest way to make dimensionless groups by comparing the hydrogen atom to another element?

the point of dimensionless groups is to understand one quantity at a time, and comparing H to another element would give you a ratio rather than an invariant/constant for hydrogen itself. plus, there is no well defined a0 for other elements...

Wasn't it said in last time's section that any true statement can be written in terms of dimensionless groups? Does this mean this statement isn't true?

It means that you cannot say anything true about hydrogen by using just those variables. Making a true (or meaningful) statement requires adding one or more variables.

I don't understand what a_0 and m_e are referring to.

Sorry, I didn't see the table on page 93.

well, it's in the text too.

where is a_0

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = \left[r^2\right] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

so is it always that if you can make a dimensionless group you have established all the relevant physics? or are there some times that you would get the wrong answer because you undimensionalized too early

So since the dimensional analysis can't be performed is that proving the model we chose was just wrong?

Yes, I think that's what we're saying here.

I'm still unclear as to what "performing dimensional analysis" means in this context. it seems to me like we're trying to find some arbitrary value without knowing why.

Does that mean only the simplest model of a problem can be dimensionless? I assumed its possible to make anything dimensionless but I guess in the context of this problem hydrogen can never be dimensionless.

I really like how in this section, some variables are chosen to use in the dimensional analysis but then it is shown that some important physics concept was forgotten and the model must be reevaluated in order to make an accurate estimation. This is really helpful because in class I was confused on how to know if you have all the right variables, and are not forgetting any.

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = [r^2] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

how do you know you have to consider relativity and quantum mechanics, and I didn't think about these at all? so is there a general approach that you can follow when you can't make a dimensionless group?

Yeah, I this didn't strike me as something to consider here. I supposed you could try to increase complexity until you reach a dimensionless group

That seems counterintuitive, although it does work here. Isn't the point of an estimation class to make things simpler? This looks like a very roundabout way to make something simple!

Simpler is relative. Compared to solving Schroedinger's equation, this approach is much simpler. But I take the global point, that maybe this example isn't the best introduction to dimensional analysis.

I've added this small example as the first use of dimensionless groups:

"As a negative example, revisit the comparison between Exxon's net worth and Nigeria's GDP. The dimensions of net worth are simply money. The dimensions of GDP are money per time. These two quantities cannot form a dimensionless group! With just these two quantities, no meanginful statements are possible."

So I understand why those cannot become dimensionless. But if you use a Force equation, then the units are the same on both sides, unlike your example of GDP and Net Worth.

what do you mean by this? you do everything relative to the speed of light?

Maybe I'm the only one, but I don't understand what relativity is? This example seems pretty complex for non-physics students and me trying to figure out the physics is taking away from my learning about dimensional analysis.

You aren't alone.

well it only took Einstein to figure it out...i'm sure one sentence is more than enough to explain..../sarcasm where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = \left[r^2\right] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

So we can simply add any dimensions we see fit in order to obtain these dimensionless quantities?

I don't understand by the speed of light intrinsically encapsulates relativity?

can we see this? i don't quite get it...

Agreed, I don't really get how it produces a size.

Well, c is a rate that has [L] and [T]. When you say size, are you referring to the [L]? This entire paragraph confuses me.

Well, c is a rate that has [L] and [T]. When you say size, are you referring to the [L]? This entire paragraph confuses me.

i agree that i still dont understand how dimensionless group=size

How do we know which size the information gives us? What tells you we can find the size for an electron and not the size of hydrogen.

Is this the radius at which the electron orbits the proton? Because wouldn't that be the size of the Hydrogen atom in this model? Unless you mean the radius of the electron itself....

I am reallyyy confused. Relativity???

where r is the distance between the proton and electron. The list of variables should include enough variables to generate this expression for the force. It could include q, ϵ_0, and r separately. But that approach is needlessly complex: The charge q is relevant only because it produces a force. So the charge appears only in the combined quantity $e^2/4\pi\epsilon_0$. A similar argument applies to ϵ_0.

Var	Dim	What
ω	T^{-1}	frequency
k	L^{-1}	wavenumber
g	LT^{-2}	gravity
h	L	depth
ρ	ML^{-3}	density
γ	MT^{-2}	surface tension

Therefore rather than listing q and ϵ_0 separately, list only $e^2/4\pi\epsilon_0$. And rather than listing r, list a_0, the common notation for the Bohr radius (the radius of ideal hydrogen). The acceleration of the electron depends on the electrostatic force, which can be constructed from $e^2/4\pi\epsilon_0$ and a_0, and on its mass m_e. So the list should also include m_e. To find the dimensions of $e^2/4\pi\epsilon_0$, use the formula for force

$$F = \frac{e^2}{4\pi\epsilon_0}\frac{1}{r^2}.$$

Then

$$\left[\frac{e^2}{4\pi\epsilon_0}\right] = [r^2] \times [F] = ML^3T^{-2}.$$

The next step is to make dimensionless groups. However, no combination of these three items is dimensionless. To see why, look at the time dimension because it appears in only one quantity, $e^2/4\pi\epsilon_0$. So that quantity cannot occur in a dimensionless group: If it did, there would be no way to get rid of the time dimensions. From the two remaining quantities, a_0 and m_e, no dimensionless group is possible.

The failure to make a dimensionless group means that hydrogen does not exist in the simple model as we have formulated it. I neglected important physics. There are two possibilities for what physics to add.

One possibility is to add relativity, encapsulated in the speed of light c. So we would add c to the list of variables. That choice produces a dimensionless group, and therefore produces a size. However, the size is not the size of hydrogen. It turns out to be the classical electron radius instead. Fortunately, you do not have to know what the classical electron radius is in order to understand why the resulting size is not the size of hydrogen. Adding relativity to the physics – or adding c to the list –

What?? This paragraph is poorly written and races through the material. What exactly are you trying to say here? Why can you arbitrarily add constants? How do you know what the effect of adding those constants is? And what is the point of this dimensionless value if we have to go to these lengths to construct it?

I, too, am curious about the cost-to-benefit ratio here. It seems that the necessary amount of physics knowledge just doubled, all to obtain a dimensionless value?

Yeah, this paragraph sort of lost me.

This confused me as well. I understand the theory behind using c, however, where does it show up in the end product?

Somehow, we need to add a constant that will allow us to remove the time dimension. The speed of light, in some way, is an available variable in all instances and will help us remove time.

I agree....especially the part where it says "produces a dimensionless group, and therefore produce a size"...I don't get this logic.

What is the classical electron radius?

It's the radius of an electron based on classical physics (rather than quantum)

I'm really hoping this section will make more sense after lecture tomorrow, because it sounds very interesting but I can't really understand it as written.

I will try!

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = \mathrm{ML}^2\mathrm{T}^{-1}.$$

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	$\mathrm{ML}^3\mathrm{T}^{-2}$	
m_e	M	electron mass
$\hbar$	$\mathrm{ML}^2\mathrm{T}^{-1}$	quantum

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = \mathrm{ML},$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)}$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)} = \text{dimensionless constant}.$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\,\text{Å} = 0.5 \cdot 10^{-10}\,\mathrm{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

Comments on page 3

I think adding more advanced information like this distracts me from the problem at hand. Maybe towards the end of the section you can mention the different ways you could approach the problem using different methods in physics?

Sorry, I'm a bit lost here. How does adding c to the list allow for radiation?

I agree, lost too....maybe a sentence of explanation?

why it allows for radiation isn't so important to this section. (I think accelerating charges give off radiation, as we learned in 8.02. Here, the electron is in circular motion, thus it has radial acceleration, so it "must" give off light. But if that happened, then it would lose some energy, and start falling inwards towards the nucleus, and it keeps radiating since it's still somewhat circular motion, and so on and so forth, until it demolishes itself in the nucleus. Obviously that's not the case, or else Hydrogen or any other atom would not be stable.)

But back to the point: Sanjoy is searching for a missing variable with a dimension of time in it so that he can use it to set up a dimensionless group.

He speculates that "c" is important in relativity, so maybe it will factor in somehow. The exact way it factors in is no important yet, since we just want to obtain something that works.

After rejecting c, because of the radiation reason, he speculates about using "h_bar", which is an important constant in quantum mechanics. How exactly it factors in isn't important for us to know here. What is important is that we don't have a reason to reject using it, and it has that dimension of time, which we were looking for.

Thus, without knowing much about the field of relativity or about quantum mechanics, you can still use constants from those fields to "arrive" at a relation, via dimensional analysis.

how do you add this by simply adding c to the list of variables?

so confused

me too

This suddenly got way too advanced. The last physics I had was 8.02

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = ML^2T^{-1}.$$

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	ML^3T^{-2}	
m_e	M	electron mass
$\hbar$	ML^2T^{-1}	quantum

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = ML,$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)}$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)} = \text{dimensionless constant}.$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\text{Å} = 0.5 \cdot 10^{-10}\,\text{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

That is an interesting point and something that was not intuitive to me at first.

While I suppose I've always considered radiation to be a form of energy loss, it didn't strike me that its effect would cause an electron to spiral into the proton.

I think that a better explanation of what radiation is and how it works would be useful here...how does radiation carrying energy away push the electron in?

I don't really understand what the purpose of this paragraph is. Is it a method that we could use, but are not going to, considering the extensive description of the next possibility?

I agree. I don't think this paragraph has actually taught me anything, about what we're supposed to be learning, it's just made me confused as to why adding in these constants apparently leads to new implications for physics.

I would recommend making a statement earlier in this reading that this section is one of those times when you have to blur your vision and not look too closely at the details of what is going on. Otherwise I think the reader might get really confused and distracted in trying to understand all the quantum physics and laws and stuff and actually end up missing the dimensional analysis lesson that you are trying to teach.

I wish I remembered Quantum better... I'm going to have to review that stuff

some of us have never taken quantum. we've been told we "major in 8.01"

Uh I'm lost, which world? The world in which the model exists?

I don't understand how you reached this conclusion. Your conclusion hinges on the fact that you arbitrarily chose to include the constant c, but couldn't you have chosen any other related constant, and had a different conclusion?

What he is saying is that if c was the proper term to add in to allow the behavior we want, electrons would actually spiral in towards the nucleus and obliterate (since quantum effects are actually what keep electrons bound to nuclei and not spiral in). So c cannot be the correct term. Planck's constant turns out to be the correct value because it is THE related constant and makes it all work. I recommend looking at a basic quantum text (such as Griffiths) for a more indepth explanation of this.

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = ML^2T^{-1}.$$

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	ML^3T^{-2}	
m_e	M	electron mass
$\hbar$	ML^2T^{-1}	quantum

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = ML,$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)}$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)} = \text{dimensionless constant}.$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\text{Å} = 0.5 \cdot 10^{-10}\,\text{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

both of these things sound very big and difficult to add

I agree - this requires more outside knowledge than I have i think.

I feel that you just guessed (intelligently) and got lucky.

drum roll.. plank's constant... Where's the wave equation?

But how does this "constant of nature" influence the physics? Before, adding the speed of light had a simple effect on our dimensional analysis, but killed the physical realities of the problem. How do we know that we are not repeating that same mistake here?

Look at any introductory quantum textbook such as Griffiths. What you find is that the angular momentum of an orbit in quantum mechanics is in integer steps of h bar.

This raises a good point, that someone asked in lecture, how do we know when we've found all of the relevant quantities? Sure h bar might be important, but why not another variable we have yet to identify, too?

Well. Physics kind of falls into categories. Classical physics is for general, everday physics. The rest of physics occurs at various 'extremes'. Relativity adds in 'c' and works with physics at extremely high speeds (approaching the speed of light, c). Statistical physics deals with large numbers of particles or interactions, another extreme. And quantum physics deals with extremely small particles and masses, and this is when 'h' becomes important. All of these physics principles are always 'true', so we could always use them, but they only become important (affect orders of magnitude) when they are in these extremes. So basically, you just need to know which extremes you are considering to know which constants you may need.

This makes quite a bit of sense. Thank you.

You might want to consider calling it the Planck constant, just so people who don't have experience with QM can know the name. Also, you should consider saying how h bar is just h/2pi, just for completeness.

How did you get the value for H?

Its a constant... not all constants are dimensionless. h is known as Planck constant

"What" might be a bit informal in this context?

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = ML^2T^{-1}.$$

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	ML^3T^{-2}	
m_e	M	electron mass
$\hbar$	ML^2T^{-1}	quantum

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = ML,$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)}.$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)} = \text{dimensionless constant}.$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\text{Å} = 0.5 \cdot 10^{-10}\,\text{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

I love your writing voice – it adds danger and heroism and intrigue (ish) to what were otherwise mundane equations in our textbooks.

Thank you!

My goal, which I'll never fully reach, is that the physical relations and quantities and properties become the actors in a story, and that we, the readers, get drawn into their story.

That's a blurry image of where I'd like to get to, and don't quite see the route. But knowing which examples and phrases go in that direction is part of figuring out the route.

Where does a_nod come from?

in the previous page, when he said instead of using a generic "r" for the coulombic force equation, we could use units of length with magnitude a_o, since it is the average distance of the electron from the nucleus in Hydrogen, aka, the "Bohr Radius"

We could put the magnitude here. It's used in the calculation below, but never printed anywhere.

I feel like this table would have been helpful on the previous page, and that things like m.e, etc. should be defined explicitly in the text.

Agreed, I was wondering about the a_0 and m_e on the last page since it references them.

This seems to have been an error: I think it was meant to be where the other table was.

Did I miss why we use electron mass and not nuclear or proton mass?

This is cool but it would have also been nice to see the final dimensionless group if you still used c as a variable.

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = ML^2T^{-1}.$$

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	ML^3T^{-2}	
m_e	M	electron mass
$\hbar$	ML^2T^{-1}	quantum

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = ML,$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e(e^2/4\pi\epsilon_0)}$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e(e^2/4\pi\epsilon_0)} = \text{dimensionless constant}.$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\text{Å} = 0.5 \cdot 10^{-10}\,\text{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

I'd like to see this also in the form of M, L, and T. I think it would help a lot with seeing how this is dimensionless to see that everything cancels out.

I agree, I'm trying to see how this is dimensionless and having a lot of difficulty.

I'd also like to see a less constants...I remember in class you talked about a system of measurement where there are no 4pi*epsilons...what was that again?

Does this mean what it did before–that we can raise this group to any arbitrary power?

I don't remember this example.

We did it in lecture and it was in the reading too. T^2*k/m was the constant there.

It's in the reading directly before this one.

This is interesting and clearly very difficult to come up with ourselves - what are your thoughts for a general approach on finding dimensionless groups?

My previous class that did dimensionless groups had a method that wasn't too bad, my guess is that he'll teach it in lecture once we get into this topic.

Read the previous few paragraphs on how this was approached.

Though I will agree, the way it's written it's not immediately clear to anyone skimming what the though process behind the equation is.

i think a simpler example to start with would have been helpful (although this was fairly fascinating). it's just a bit hard to understand a concept when the roadblocks are brought up before we ever see the first solution

I agree. This example seems pretty complex and thus loses it application purposes. Maybe a simpler example?

The thing is, when talking about atomic physics, this _is_ the simplest example.

Yea I agree, I think the only way I might have gotten this on my own is by backtracking the units to make sure the entire thing is dimensionless at the end.

So once we introduce h, then we can write the above expression as dimensionless...so it is still true

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = \mathrm{ML}^2\mathrm{T}^{-1}.$$

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	$\mathrm{ML}^3\mathrm{T}^{-2}$	
m_e	M	electron mass
$\hbar$	$\mathrm{ML}^2\mathrm{T}^{-1}$	quantum

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = \mathrm{ML},$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)}$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)} = \text{dimensionless constant}.$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\,\text{Å} = 0.5 \cdot 10^{-10}\,\mathrm{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

isn't there a constant multiplying this value?

Could you also have a table that included these constants somewhere on this page? It'd be nice to see them and know what they represent.

That's pretty neat.

this is using constants which were derived in the study of the atom though? isn't it like going around in a large circle?

But it's taking physical values and making them dimensionless. You still don't know if there was some non dimensional constant to begin with (i.e. K=200 or something)

We actually did this in 8.04 to familiarize ourselves with dimensional analysis. We had a whole unit on it.

What you you mean by "it turns out to be 1" How did you calculate this?

It means that that the answer for ao is very close to the true answer, so we're not missing some dimensionless constant of 2pi or G or 1 billion

ha that always impresses me

Overall, this section was probably one of the most technical reads in the course thus far, but after reading it a second time I understand the point you were trying to make. Still, however, I think the large amount of technical-heavy material in this section makes it a little intimidating.

I definitely agree. This section definitely required me to read over it a couple of times before I understood. I like how you vary between more technical examples and more everyday examples though.

I think highlighting the process a little more helps. All the text is hard to sift through.

I agree with these points as well, it took me far longer to go through these short paragraphs because of jumps in concepts, although only later did I realize the main point of dimensional analysis given what you know.

Why is this important then? This doesn't seem to give us any new information

allows radiation. So the orbiting, accelerating electron would radiate. As radiation carries energy away from the electron, it spirals into the proton, meaning that in this world hydrogen does not exist, nor do other atoms.

The other possibility is to add quantum mechanics, which was developed to solve fundamental problems like the existence of matter. The physics of quantum mechanics is complicated, but its effect on dimensional analyses is simple: It contributes a new constant of nature $\hbar$ whose dimensions are those of angular momentum. Angular momentum is mvr, so

$$[\hbar] = \mathrm{ML^2T^{-1}}.$$

The $\hbar$ might save the day. There are now two quantities containing time dimensions. Since $e^2/4\pi\epsilon_0$ has T^{-2} and $\hbar$ has T^{-1}, the ratio $\hbar^2/(e^2/4\pi\epsilon_0)$ contains no time dimensions. Since

Var	Dim	What
a_0	L	size
$e^2/4\pi\epsilon_0$	$\mathrm{ML^3T^{-2}}$	
m_e	M	electron mass
$\hbar$	$\mathrm{ML^2T^{-1}}$	quantum

$$\left[\frac{\hbar^2}{e^2/4\pi\epsilon_0}\right] = \mathrm{ML},$$

a dimensionless group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)}.$$

It turns out that all dimensionless groups can be formed from this group. So, as in the spring–mass example, the only possible true statement involving this group is

$$\frac{\hbar^2}{a_0 m_e (e^2/4\pi\epsilon_0)} = \text{dimensionless constant}.$$

Therefore, the size of hydrogen is

$$a_0 \sim \frac{\hbar^2}{m_e (e^2/4\pi\epsilon_0)}.$$

Putting in values for the constants gives

$$a_0 \sim 0.5\text{Å} = 0.5 \cdot 10^{-10}\,\mathrm{m}.$$

It turns out that the missing dimensionless constant is 1: Dimensional analysis has given the exact answer.

Pretty cool!

sweeeet. I like it when the approximations are exact

agreed, but i also like when i can reproduce them easily!

it is easy!

This was another really good example. I think this section has been the best one so far, and possibly the most useful (for me).

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1 Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3 Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

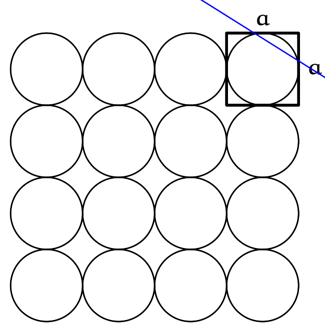

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3\,\text{Å}$, and has mass Am_{proton}. The density of the substance is

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of N_A/N_A, where N_A is Avogadro's number:

$$\rho \sim \frac{Am_{\text{proton}}N_A}{(3\,\text{Å})^3 \times N_A}.$$

The numerator is A g, because that is how N_A is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_A. However, N_A is more familiar than m_{proton} because N_A arises in chemistry and physics. Using N_A also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g}\,\text{cm}^{-3}.$$

Is this an exact statement? I don't mean to nit-pick, but the fact that you don't say about or equivalent could confuse someone as to whether it is exact, a definition, or simply an approximation.

1A = 0.1 nanometers (which equals 1*10^-10 meters)

Yes, it does. But he explained this above where he said the radius was .5 Angstroms. Which he multiplied by 2 to get the diameter.

do you mean "of the rule of thumb"? otherwise, I'm confused.

Why is this a consequence of hydrogen having a diameter of 1A? I'd appreciate it if someone could please elaborate for me.

I agree. I have no idea why Hydrogen having a diameter of 1 Angstrom means that the typical interatomic spacing is on the order of 3 Angstroms. Also, your sentence has too many copies of the word "is" in it.

Consequence is probably a poor word choice for it but it means that from that we can guess what the typical diameter might be, and since hydrogen is the smallest is gets, 1 A would be the lower bound (there's a later note that explains this). However, this could be explained a bit better (I was confused until I read the later note... I thought it meant that you have hydrogen with a 1 A diameter and in additional to that there is the space of 3 A between it and the next hydrogen... which doesn't make sense.)

Can we justify the 3 angstrom interatomic spacing? Assuming that the van Der Waals radius is 0.5 angstroms does not yield 3 angstrom separation in an obvious way. In the last section, you picked 1 as the dimensionless quantity. Taking a hint from quantization, do you just pick 2 for larger atoms and average out to 1.5?

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3\,\text{Å}$, and has mass Am_{proton}. The density of the substance is

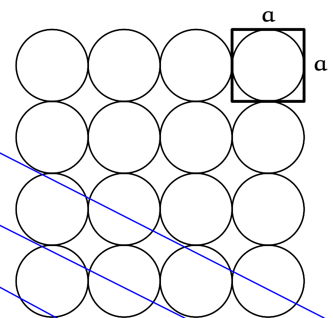

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of N_A/N_A, where N_A is Avogadro's number:

$$\rho \sim \frac{Am_{\text{proton}}N_A}{(3\,\text{Å})^3 \times N_A}.$$

The numerator is A g, because that is how N_A is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_A. However, N_A is more familiar than m_{proton} because N_A arises in chemistry and physics. Using N_A also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g}\,\text{cm}^{-3}.$$

How did we get 3 from 1?

The 1 was the diameter of hydrogen. The 3 is typical interatomic spacing.

The 3 and 1 represent 2 different values- 3A is the typical space between all atoms while 1A is the diameter for Hydrogen specifically.

So hydrogen atoms are spaced by only 1A when most other atoms are typically spaced between 3A? Am I following you correctly?

Why doesn't this vary significantly based on the size of the atom? Or is three just the overall average.

As you say, it is based on the size of the atom. Hydrogen is at the very small end, and uranium is at the large end. 3 Angstroms is a good average size to use for the common atoms in ordinary substances. As a *very* rough approximation, think of the diameter as 1 Angstrom per shell. (The number of shells is the row number in the periodic table.)

Too many 'is's in this sentence.

This paragraph's wording is a little bit confusing.

I think the confusion also comes from introducing the diameter size earlier with the unit Angstroms.

I realize A is commonly used for atomic mass, but maybe in this particular example, where you just introduced angstroms and are talking about atoms with diameters equal to a.... using M for mass might be a little less confusing.

or maybe even 'N' for number of protons/neutrons. Just not 'A'

I saw this too...why can't we use M instead? Armstrong and A for mass is a bit confusing.

you can't use M because the number is not actually a mass. I wasn't confused by Angstroms & A...I think they are different enough to work. if you want something other than 'A', 'N' would probably be the best option.

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3$ Å, and has mass Am_{proton}. The density of the substance is

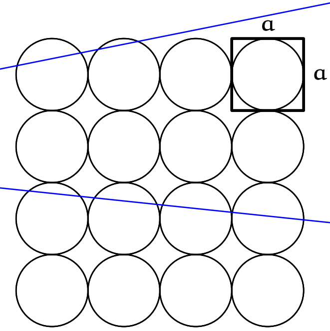

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{proton}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of N_A/N_A, where N_A is Avogadro's number:

$$\rho \sim \frac{Am_{proton}N_A}{(3\,\text{Å})^3 \times N_A}.$$

The numerator is A g, because that is how N_A is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_A. However, N_A is more familiar than m_{proton} because N_A arises in chemistry and physics. Using N_A also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g cm}^{-3}.$$

If it's not a mass, don't call it a mass. Just introduce it as the number of protons and neutrons (more or less).

I disagree, I think that most people are familiar with the concept of atomic mass.

If A is a mass then it must have units of mass. Otherwise it is not a mass.

I agree. This sentence makes no sense.

A is an atomic mass

Saying A is the atomic mass, then saying its not, then multiplying it by m_proton is a little confusing to follow at first. Would it make more sense to simply say A for an atom is roughly the number of protons and neutrons? This way it would make more sense to simply multiply by the mass. I understand why it is explained this way...it makes perfect sense but I feel it could be explained a little bit simpler and easier for someone to follow who isn't familiar with atomic properties (although I'm sure most all of this would then be overwhelming).

I think this makes sense... but maybe that's just because I'm not unfamiliar with the term "atomic mass".

I found the use of 'A' for atomic mass and 'a' for side length a bit distracting at first...the first time i read it my brain didn't want to parse it correctly.

Let a A be the "atomic mass" of a particular atom, roughly equal to the number of protons and neutrons in the nucleus of the atom. Although A is called a mass, it is actually dimensionless; in reality, an atom has a mass of Am(proton). Each atom occupies a cube of side length a 3Å, making the density of a substance:

are we saying this is dimensionless (i.e.) as part of a dimensionless group or it is it just dimensionless by itself

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

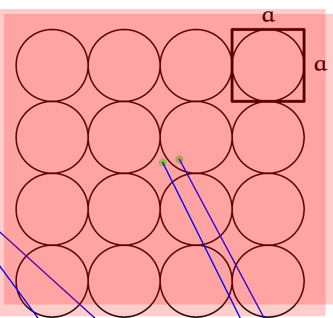

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3\,\text{Å}$, and has mass Am_{proton}. The density of the substance is

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of N_A/N_A, where N_A is Avogadro's number:

$$\rho \sim \frac{Am_{\text{proton}}N_A}{(3\,\text{Å})^3 \times N_A}.$$

The numerator is A g, because that is how N_A is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_A. However, N_A is more familiar than m_{proton} because N_A arises in chemistry and physics. Using N_A also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g}\,\text{cm}^{-3}.$$

Why is it dimensionless?

It appears that it is more of a count (the number of protons and neutrons), and when the unit of mass is needed, A is multiplied by the mass of a proton.

It;s more of just a number. Since we approximate the weight of neutrons and protons to be the same, it's more of just giving a number of how many. This isn't entirely true in my opinion, since that's the definition of 'amu' or atomic mass unit.

If this is true, then its confusing in the paragraph that it says to let A be teh atomic mass of the "atom". If a is just a count, then shouldn't A be considered to be something like Avogadro's number?

I believe it's just convention to call it "mass".

the mass is approximately equal to the number of protons and neutrons, so it's just a count. but my question is, if this really is a mass, it must be convertible into other mass units like grams. if this is true, it can't be dimensionless anymore right?

Although 'atomic mass' is often given in AMU, I think, which are actually units of mass.

Is it still valuable while dimensionless or is it only useful after converting it back to the dimensioned version?

It comes up on this page but the idea is that you are using the mass of the proton also so you can just use A as a completely dimensionless variable that depends on the element.

Is m proton the mass of a proton or is it the mass of the protons in the atom? I am assuming its the mass of a proton but I just wanted to clarify.

I think this could be much smaller and still get the point across.

I feel like I have seen this drawing before...where it asks you what is the area of wholes in between? I don't know if this is relevant but it is still interesting.

It looks like you mean A is a scaling factor on the mass of a proton, and not that A is actually a mass itself.

If this is true, I would recommend using a different letter since the A might be confused with the angstrom symbol.

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3\,\text{Å}$, and has mass Am_{proton}. The density of the substance is

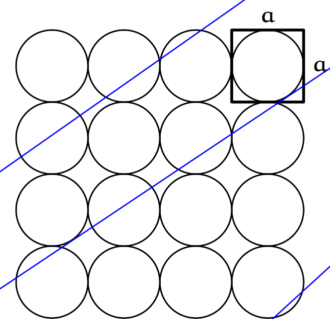

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of N_A/N_A, where N_A is Avogadro's number:

$$\rho \sim \frac{Am_{\text{proton}} N_A}{(3\,\text{Å})^3 \times N_A}.$$

The numerator is A g, because that is how N_A is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_A. However, N_A is more familiar than m_{proton} because N_A arises in chemistry and physics. Using N_A also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g}\,\text{cm}^{-3}.$$

Don't you still have to remember Avogrado's number now?

Oh sorry, you say that below.

I like this. Despite being a little unclear at first, this is a really neat trick to move on in the problem.

I am entirely confused on how this does anything.

Although you mention what NA is a few lines below, it might be useful to just include its value here in parentheses.

He defines this value when he uses it about 2 lines further down. I think that this is sufficient.

I agree with 10:14. We only really care about the numerical value of N_A when we plug in numbers.

Personally, I would prefer that he would state a value immediately after he defines/discusses a new variable.

So what does a mole of protons weight?

It's almost exactly 1 gram. One mole of carbon-12 atoms weight exactly 12 grams, and carbon-12 has the almost the same mass as 12 protons (the 6 neutrons and 6 protons add up to roughly 12 protons, and the 6 electrons are very light).

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3\,\text{Å}$, and has mass Am_{proton}. The density of the substance is

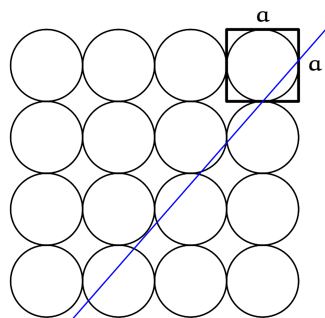

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of N_A/N_A, where N_A is Avogadro's number:

$$\rho \sim \frac{Am_{\text{proton}}N_A}{(3\,\text{Å})^3 \times N_A}.$$

The numerator is A g, because that is how N_A is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_A. However, N_A is more familiar than m_{proton} because N_A arises in chemistry and physics. Using N_A also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g}\,\text{cm}^{-3}.$$

I had to read this sentence a couple times to really understand it, I don't know if it should be phrased better or if I just personally couldn't wrap my head around it.

i got confused because i thought g = gravity but i think it's grams here. it'd probably make sense to say Na * mproton = 1g so we understand that this is why it turns to A g.

This is because a mole of hydrogen atoms just weighs 1 gram as given by the periodic table, correct?

Yeah I feel like it would make it a little more clear if gram is actually written out instead of writing g.

Absolutely thought that was a 'g' for acceleration. Perhaps put it in parenthesis like (in grams)

me too, at first, I thought he meant gravity.

actually, I've done this on the TI-89 too, if you type just g for gram, it will give you the gravitational acceleration instead of gram. "gm" is the symbol used for grams on the calculator.

I'd like to be reminded in the text what Avogadro's number is, with units.

No estimating here?

Why do you need to estimate to multiple 3 by 6?

no units here? should be cm^3

We only estimated when multiplying for large numbers and powers of 10. Basic arithmetic would be unnecessary estimation

if it is easier to not estimate, you should give the accurate answer

why is there a "however" starting this sentence? I think you can just get rid of the however.

The however makes sense because at first it seems like he's contradicting himself a little but this sentence explains how he's not.

So in the end this doesn't save us a constant, but allows us to use a more familiar/memorable constant?

add: "frequently"

5.3.2 Atomic sizes and substance densities

Hydrogen has a diameter of 1Å. A useful consequence is the rule of thumb is that a typical interatomic spacing is 3Å. This approximation gives a reasonable approximation for the densities of substances, as this section explains.

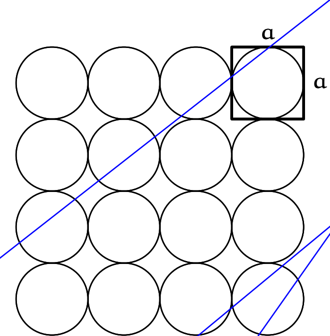

Let A be the atomic mass of the atom; it is (roughly) the number of protons and neutrons in the nucleus. Although A is called a mass, it is dimensionless. Each atom occupies a cube of side length $a \sim 3\,\text{Å}$, and has mass Am_{proton}. The density of the substance is

$$\rho = \frac{\text{mass}}{\text{volume}} \sim \frac{Am_{\text{proton}}}{(3\,\text{Å})^3}.$$

You do not need to remember or look up m_{proton} if you multiply this fraction by unity in the form of N_A/N_A, where N_A is Avogadro's number:

$$\rho \sim \frac{Am_{\text{proton}}N_A}{(3\,\text{Å})^3 \times N_A}.$$

The numerator is A g, because that is how N_A is defined. The denominator is

$$3 \cdot 10^{-23}\,\text{cm}^3 \times 6 \cdot 10^{23} = 18.$$

So instead of remembering m_{proton}, you need to remember N_A. However, N_A is more familiar than m_{proton} because N_A arises in chemistry and physics. Using N_A also emphasizes the connection between microscopic and macroscopic values. Carrying out the calculations:

$$\rho \sim \frac{A}{18}\,\text{g}\,\text{cm}^{-3}.$$

This section is great! Back in the days when I competed in science competitions, I would often use dimensional analysis to guess the answer when I did not understand the physics. It helped me when in a few cases!

I like this jump, it really shows how number's like Avogadro's number that we learned back in high school can pop up and still be useful

This conclusion is really cool and makes a lot of sense in the approximation world. It reminds me of the period tables with size of atoms plotted

I also really like when these methods that I would never think of eventually reach an equation that makes sense to me. It's also nice to see in the table below that it's a good approximation.

95

The table compares the estimate against reality. Most everyday elements have atomic masses between 15 and 150, so the density estimate explains why most densities lie between 1 and 10 $g\,cm^{-3}$. It also shows why, for materials physics, cgs units are more convenient than SI units are. A typical cgs density of a solid is 3 $g\,cm^{-3}$, and 3 is a modest number and easy to remember and work with. However, a typical SI density of a solid 3000 $kg\,m^{-3}$. Numbers such as 3000 are unwieldy. Each time you use it, you have to think, 'How many powers of ten were there again?' So the table tabulates densities using the cgs units of $g\,cm^{-3}$. I even threw a joker into the pack – water is not an element! – but the density estimate is amazingly accurate.

Element	$\rho_{estimated}$	ρ_{actual}
Li	0.39	0.54
H_2O	1.0	1.0
Si	1.56	2.4
Fe	3.11	7.9
Hg	11.2	13.5
Au	10.9	19.3
U	13.3	18.7

Comments on page 5

Your lengthier readings usually have sections or subparts and a lost of forecasting of what is to come. This passage seemed to read more like a physics textbook and thus was harder to follow, at least for me.

A bit too high with respect to the other line its on, formatting issue

i feel like it would be useful to include an example or two of what these densities mean in real-life terms, so that we can get a feel what these values mean.

All the estimates seem to be 10^0.5 off that seems kind of significant based on the estimatation we have done in this class.

Well, I don't think we would worry about a power of 2 here and there, but it does seem that we have a consistent underestimate of rho, which does suggest that we've missed something. Fortunately, whatever it is isn't too important.

this is very true i think the units in cgs are more intuitive or we are more familiar with

cgs = ??

cgs = centimeters, grams, seconds.

Why don't you just use cgs throughout this section? (as suggested on the first page, it would also simplify the force equation)

I think it depends on your background for which units you are more familiar with. As a MechE, I know m/kg/s much better (although not in a Physics context).

Just be thankful we aren't using miili-inch - Slug -pound system.

When I first wrote (don't ask how long ago), it was in cgs. But I found that almost all teaching, even for electromagnetism, is in SI units (meters, kilograms, seconds). That was especially true in England, but it's true here too. So I thought that the correct choice for the long term is to use SI units even though they are, for electromagnetism, messier than cgs units. (In cgs units, the 4*pi*epsilon_0 is replaced by 1.)

The table compares the estimate against reality. Most everyday elements have atomic masses between 15 and 150, so the density estimate explains why most densities lie between 1 and 10 $g\,cm^{-3}$. It also shows why, for materials physics, cgs units are more convenient than SI units are. A typical cgs density of a solid is 3 $g\,cm^{-3}$, and 3 is a modest number and easy to remember and work with. However, a typical SI density of a solid 3000 $kg\,m^{-3}$. Numbers such as 3000 are unwieldy. Each time you use it, you have to think, 'How many powers of ten were there again?' So the table tabulates densities using the cgs units of $g\,cm^{-3}$. I even threw a joker into the pack – water is not an element! – but the density estimate is amazingly accurate.

Element	$\rho_{estimated}$	ρ_{actual}
Li	0.39	0.54
H_2O	1.0	1.0
Si	1.56	2.4
Fe	3.11	7.9
Hg	11.2	13.5
Au	10.9	19.3
U	13.3	18.7

I don't know if the typical reader will be acquainted with cgs units. Would it be useful to include a brief note on what they are, or should we just assume the reader can look it up if he is confused?

I didn't know and I had to look it up

Yea, I feel that even mentioning in parenthesis that cgs is centimeter, gram, seconds would be useful. Solid Works is the only reason why I recognized this.

Thanks for the clarification - I had no idea what cgs stood for. Consequently, I agree this could use a little explanation in the text.

I wasn't familiar with cgs either. We always used kilograms/meters/seconds in physics.

Woah...who would have thought. Why are some of these extremely close and others very off?

If I remember correctly there is some interaction with the way orbitals fill up that causes atom sizes to shrink across a period even though each the size increases going down a group (http://en.wikipedia.org/wiki/Atomic_radius#Calculated_atomic_radii)

And the density will always be too small, as we are approximating an atom as a box, and the volume is in the denominator. And so, if you have something spaced smaller than average, then the volume will be really way small.

Yeah, all the estimates are lower than the actual.. Can we do something about this?

It's the opposite effect to what we found when estimating the maximum cycling speed. There, v was proportional to power^(1/3), so even large errors in estimating power turned into small errors in estimating v.

Here, however, we are estimating density, which depends on diameter^3. So even small errors in estimating the diameter (i.e. small deviations from the 0.3 nm baseline value) produce a large change in the density.

Is there a reason that this one is the worst? Does it have anything to do with iron's magnetism?

probably a propagation in error of interatomic spacing

The table compares the estimate against reality. Most everyday elements have atomic masses between 15 and 150, so the density estimate explains why most densities lie between 1 and 10 $g\,cm^{-3}$. It also shows why, for materials physics, cgs units are more convenient than SI units are. A typical cgs density of a solid is 3 $g\,cm^{-3}$, and 3 is a modest number and easy to remember and work with. However, a typical SI density of a solid 3000 $kg\,m^{-3}$. Numbers such as 3000 are unwieldy. Each time you use it, you have to think, 'How many powers of ten were there again?' So the table tabulates densities using the cgs units of $g\,cm^{-3}$. I even threw a joker into the pack – water is not an element! – but the density estimate is amazingly accurate.

Element	$\rho_{\text{estimated}}$	ρ_{actual}
Li	0.39	0.54
H_2O	1.0	1.0
Si	1.56	2.4
Fe	3.11	7.9
Hg	11.2	13.5
Au	10.9	19.3
U	13.3	18.7

...of a solid "is" 3000 kg m^-3.

I was wondering if we were going to have to make special considerations for molecules instead of elements when we were deriving the equation before but I guess it makes sense that we don't need to. We never needed to assume it was just one element.

It would make more sense if it held for only one element! I wonder if CO2 is as close as H2O...

I don't see how this could possibly apply to gases. Typical gas densities are 10^-3 g/cm^3. So A would have to be a small fraction of 1.

I assumed that these densities were talking about solids until water came up (I realize water is a weird case though with regards to liquid/solid densities). Is it fair to assume that the liquid and the solid densities of most of these elements are close enough?

even though water is polar? does this make it less than ideal?

I was wondering about this as well

I wasn't thinking about that, but now I am. putting myself on the thread so I will get an email when someone answers the question for us.

Being polar doesn't affect the intermolecular spacing much in water. The molecules are still as closely packed as possible (basically, until their electron clouds touch and repel each other).

But in ice, being polar is responsible for the open structures that water molecules form – which is why ice is less dense than water (unlike most liquids, which contract when they freeze). So for ice, the polarity does slightly affect the average spacing.

The estimate for water is actually absurdly accurate, much closer than any of the others, even though it was not what we were modeling. Interesting.

The table compares the estimate against reality. Most everyday elements have atomic masses between 15 and 150, so the density estimate explains why most densities lie between 1 and 10 g cm^{-3}. It also shows why, for materials physics, cgs units are more convenient than SI units are. A typical cgs density of a solid is 3 g cm^{-3}, and 3 is a modest number and easy to remember and work with. However, a typical SI density of a solid 3000 kg m^{-3}. Numbers such as 3000 are unwieldy. Each time you use it, you have to think, 'How many powers of ten were there again?' So the table tabulates densities using the cgs units of g cm^{-3}. I even threw a joker into the pack – water is not an element! – but the density estimate is amazingly accurate.

Element	$\rho_{\text{estimated}}$	ρ_{actual}
Li	0.39	0.54
H_2O	1.0	1.0
Si	1.56	2.4
Fe	3.11	7.9
Hg	11.2	13.5
Au	10.9	19.3
U	13.3	18.7

I have some problems with this section. Firstly, it was difficult for me to understand, but when go back to reread it, I don't find Topic sentences, or even a decent conclusion that summarizes the Section. I understand this is hard material to teach, but you can't expect the reader to reread and reread your sections until they make sense. You should reiterate the points of interest yourself.

Agreed. Some basic introductions and conclusions, like you have in the previous sections, would really help make this section clearer.

Agreed!

In the beginning of 2.006 Prof. Brisson taught dimensional analysis using the MLT quantities that you described here. He had a succinct method that was very easy to understand and was applicable to the questions of scaling we saw before. I think you should talk to him and get his notes on it, because this section was entirely too confusing for me, and I know how to do it!

Thanks for the suggestion. I've just asked him.

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

This was definitely one of my favorite chapters so far, I've always been incredibly interested by this material but as a course 2 don't get to see it too much. I keep being impressed with the broad spectrum of material we cover that is relevant to everyday life

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

Read this section for the memo due on Sunday at the end of spring break.

this is an interesting question to start the chapter

I'm confused by this - I thought light did feel the effect of gravity, especially when you consider phenomena like black holes? Or is that something different entirely?

I agree, interesting questions. I mean I guess you could argue that you feel light as energy, but not so much as something like drag. It doesn't retard any force

I was confused by the wording of this at first as well... I don't think the question is stating that light doesn't feel the effect of gravity, its more like "how does gravity affect light?"

Light is affected by gravity. But since we normally think of it as "massless photons" we assume that something that acts on mass, ie: gravity, will not affect it. This is an incorrect assumption.

so does this mean that everything is affected by gravity? I mean I thought light was a form of electrog-magnetic radiation which is a type of energy. So does this mean that energy is affected by gravity (even magnetic energy or electric energy)?

I've never even thought of that question.

I seem to think about this question a lot, most importantly when it refers to stars. I don't think we would see stars at all if light were affected by gravity.

This is such a great question to begin the section with. It attracts my attention because it is something that I have never ever thought of.

I agree - its also a question that the reader can start to try to reason about...what do we know about black holes? do we ever hear about light bending?

It's already very attention-grabbing, but maybe even mentioning black holes in the reading would be cool.

I agree about black holes, or maybe just mention something in class about them. I'm pretty interested in them and think it would be nice to go into a bit of detail since they are related to this topic

I've wondered this myself.

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

I am kind of confused by this... the first part says light is not effected and the second part says we will analyze how it is effected...

It's a writing / teaching technique to say something that most people will agree with, then prove them wrong.

Is this similar to how the cloaking/invisibility technology theoretically works?

Could someone give some intuition into how gravity bends light?

Ever hear of a black Hole? Black holes suck in light because the mass is so dense that the escape velocity is greater than the speed of light. If you understand this concept, it may be easier for you to justify deflection from gravity.

I'm not too versed in physics, but the way I understand it at a subatomic level, mass and energy and mass can be thought of as interchangeable values (remember E=mc^2). Anything made up of energy (which is everything) has a mass. Please correct me if I am wrong.

So gravity bends light because, it turns out that mass just bends "space time" into different shapes, and both massive and massless particles follow geodesics, or the shortest path between two points. Since gravity in the general relativistic picture isn't just a force between two masses, but instead is just things moving in "straight lines," light bends just like anything else.

So are we going to calculate how gravity bends light, in order to see that it doesn't affect light unless the gravitational force is extremely strong? Because I know a black hole does this.

Why has this concept never been introduced in Physics classes?

I'm really interested in the applications to black holes.

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

I've never even heard of these!

Neither had I! According to Wikipedia, a geodesic equation in math is one that generalizes the notion of a straight line to curved spaces. The Wikipedia articles has a great figure to refer to and a more detailed explanation.

Ah, thanks for the explanation. Saved me a Google search.

I was wondering if these were related to geodesic domes (think of Spaceship Earth at Epcot). The wikipedia article on geodesic domes is also really interesting.

I think putting a little more explanation leading into this equation in the reading would be helpful

I get the impression that my "huh? what?" thoughts for this paragraph will make the analysis more awesome.

It might be helpful to explain a few terms here. "geodesic" and "metric tensor" are completely new terms for me, allowing me to get lost in just the 2nd sentence.

I think that's the point. It's supposed to demonstrate how confusing and complex the situation is.

I agree that definitions might be helpful. However, the point is not to use the equations but to show how we can get around them. In that sense, definition is not absolutely necessary.

I understand that these equations are generally listed to demonstrate how confusing certain problems are, but I don't know how necessary this is to do every time (we had another example before, either with flow or pressure, can't remember exactly).

This is pretty cool stuff that you never learn in engineering courses (at least at MIT). It think it would be nice to explain this stuff, even if it is not essential to the problem we are trying to solve.

No, because then we get complaints about how confusing and unnecessary the definitions of equations and terms of throw-away examples are. This way, people can just look up terms they feel impede their understanding without messing with the flow.

What about the beta mu and nu?

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

pretty sure I've never seen these before

Similar to notes posted elsewhere: I don't think many have, and that is the point. They are supposed to look "mean and scary" so as to make our, hopefully easier, method look nice.

Yeah, but this one is a new level of scariness.

Coming at the end of spring break, this looks quite frightening!

Yeah..I don't even understand the formula..

spelling error

you repeat the "whose evaluation require solving" twice...maybe phrase it differently

i believe it's done to emphasize the fact that it will be difficult to solve this equation.

what is the difference between a metric tensor and say the tensor used in structures?

If the intent of this is to repel us from solving it traditionally/accurately, it is succeeding admirably.

I believe it is and you are most correct.

My takeaway from this paragraph was to be prepared to cry if you have to solve these for any non-trivial case

Indeed! There are very few cases where the solutions are known. One of the few cases is for a spherically symmetric, nonrotating mass – the Schwarzschild solution (derived while Schwarzschild was in the German army in WW1). It was such a feat of equation solving that even Einstein did not expect an exact solution.

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

if the point of this paragraph was to say that these equations are normally very difficult and hard to understand, then job well done. :)

hahaha agreed

Indeed, though to analyze this question in terms of syntax, the double use of phrases that start with 'whose' seems awkward.

I agree. Mentioning Christoffel symbols and curvature equations is unnecessary and detracts from your point, which is that these equations are hard to solve, and quite likely not something the reader even knows how to approach.

I agree - you use this same technique of telling us how complicated solving out these giant equations is in previous sections. I think it is a tactic that only really needs to be included once because realizing how complicated brute force is isn't really specific to each equation - it is more of a general concept that we learned in a previous section. I think most people can recognize its application after just looking at each new jumbled mess of symbols.

Giving a description in really complicated terms kind of makes me not want to read the section though. I like the point you're making, but I like the method of saying "solving this would really suck" and then going into a better way to approach the problem. Like you do in the next paragraph.

I think this paragraph and the beginning section makes me want to skim. I'm wide-eyed from reading it and it just feels cumbersome. I feel like it could be simplified quite a bit, and maybe a side note to an appendix for more information (for the actual book). For example, "There are really difficult calculations here, but we can use the familiar principle..." A bit more than that definitely, but the way it is just feels hard to read.

i think that showing this complicated equation and talking about the ten non-linear partial-differential equations is a great way to let the reader understand why the approximation we are about to do is so powerful and how it makes life so much easier!

yeah! this paragraph is scary! I think terms like metric tensor, relativity curvature, geodesic, ten non linear partial differential equations really brings out how tough this problem is to solve actually. to top that off, some frankensteinish name like Christoffel really tops it off.

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

Ten partial-differential equations! Thank god I am not a physicist.

....And they aren't even linear equations.

I'm a little confused about how it could be 10. Normally these kinds of things represent 3 or 9 because there's 3 directions which can be represented with respect to the same 3 directions. I definitely don't know enough to actually think the 10 is wrong, it just surprises me.

So there are ten if you're talking about 3 space dimensions and one time dimension because it is not a vector equation, it's dealing with tensor-like things. The reason it is ten is more involved... but that's why it's not three.

Agreed, this equation sounds insane and double agreed on the physicist comment.

Without symmetry, it would be 16 equations: 4 possibilities for the first index (mu) times 4 possibilities for the second index (nu). Then symmetry reduces the number from 16 to 10. The symmetry operation is switching index 1 and 2, and it doesn't change anything physical, i.e. it doesn't give you a new equation.

This wording seems off to me. Maybe "The equations are numerous, and worse, nonlinear." I'm not really sure (I'm not a literary genius), but it does not quite read right to me.

what;s gravity

so this is explaining the "brute force" approach to solving the problem?

what about numerical approximations like the Runga-Kutta method?

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

My new reply for when my parents ask what I'm learning in school.

but don't you think that's a weird thing to put in a textbook? i wouldn't say we're lowering standards, i would say we're looking for a different answer.

I read it as kind of a joke. It does sound a little weird for a textbook if taken out of context but here I think its clear that we are lowering the expectations of accuracy in the interest of saving ourselves from spending 10 years learning the math behind it.

I agree. Its a trade, we get decreased accuracy, but it can actually be solved in a class period

I like how this ties back to a previous chapter- this book flows really nicely

I agree but I don't think it would be a bad idea to mention what former chapter it does refer to.

that still seems kinda high

I think you're off by a factor of 10 or so.

Yeah this is strange...one thousand pages? I thought we were supposed to make the problem significantly easier by sacrificing some accuracy?

Misner-Thorne-Wheeler, the great bible of GR is 1215 pages. Maybe we are just measuring against that.

yeah i think the idea is to say that it'll be fewer than the "10 years, 1000 pages, 10 non-linear diff eqns." in this case, it will be FAR fewer.

I didn't catch the reference to this great bible of GR, thanks for the note.

I didn't get that either. Thanks

What is that?

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

There are 4 steps listed below.

haha. Nice catch. Completely missed that.

In another note, have these "three/four steps" been mentioned before in this way? It's been a week since I've read a memo but I don't remember seeing them spelled out this clearly before. If not, maybe this should be in an earlier section.

I don't think they've been numbered like this but we have used these steps before

I think its a nice little summary of what we've been doing.

I think enumerating the steps in this way is very useful and should be added when dimensional analysis is introduced.

Of the four steps, I still feel like this one is the most difficult. It requires a lot of initial thinking, whereas manipulation of the dimensionless groups seems less thought-intensive.

I can already see how this is going to be more useful than it was in the last unit

So we're not using the equations at all?

Well, we are, just not that particularly nasty one.

Right, I think we're going to form dimensionless groups which will show us the important ratios we should care about.

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a night-mare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

i like the fact that you listed out the steps...aside from the fact that there are 4 listed steps here and you wrote "three"

I agree. Steps are necessary for any method and I don't believe they were enunciated before.

Yeah, these steps would be nice to see in the previous chapter (the one that introduces dimensionless analysis)

I agree that the steps help. I usually list steps when learning a new method for approaching a problem. Could you list steps for other methods? I think it would help to understand them.

Yeah I really like how the steps summarize what we will be doing in this chapter. It makes it easier to follow. Also, this provides a start to solving problems that may seem daunting, and I think using this will raise my confidence when dealing with such problems.

I, also, really like the fact that you listed the steps here...i think that doing something similar would work really well in a couple of the more complicated examples from earlier.

how does one define "most general" here? i'm not sure what that means in math terms.

Before, we've tried to incorporate as many terms as possible, so perhaps it's something involving that. Then 'most general' would mean not neglecting a parameter that shows up in the dimensionless groups.

most general basically means describing the phenomena using the most simple and fundamental variables (which are arranged into dimensionless groups)

How do we know what the most general dimensionless statements are?

By "most general" I mean the form in which one dimensionless group is a function of all the other dimensionless groups. That's as general a statement as you can make. (Then you add physical knowledge to restrict the statement.)

Several of you suggested that I include the four-step (not three-step!) recipe in earlier sections. That's a good idea, and it would also make it clearer here what I mean by "most general" dimensionless statement.

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

If possible, I think it would be nice to include more explanation about this point before this reading. I believe in the previous reading, this corresponds to the section where you debated including c or h in the dimensionless value, and I was very confused where all of that was coming from. Perhaps a reordering of sections would make these points more clear.

Is this kind of like divide and conquer? using items you already know to assist in making the dimensional analysis of the unknown.

I think you have to take a little from divide and conquer for just about every problem you ever do, however this is a more detailed approach to answering questions.

great process.

"are followed" or "are covered"?

"We use these steps in the following sections"? ..."are done" does not read right.

I don't think this last comment is that necessary.

I think it actually guides the reader.

I think you should say somewhere above that the thing you are going to figure out is the angle.

I'm kinda having trouble with the idea of gravity bending a photon. Could you maybe put a little in about that concept?

5.4 Bending of light by gravity

Rocks, birds, and people feel the effect of gravity. So why not light? The analysis of that question is a triumph of Einstein's theory of general relativity. We could calculate how gravity bends light by solving the so-called geodesic equations from general relativity:

$$\frac{d^2x^\beta}{d\lambda^2} + \Gamma^\beta_{\mu\nu}\frac{dx^\mu}{d\lambda}\frac{dx^\nu}{d\lambda} = 0,$$

wwhere $\Gamma^\beta_{\mu\nu}$ are the Christoffel symbols, whose evaluation requires solving for the metric tensor $g_{\mu\nu}$, whose evaluation requires solving the general-relativity curvature equations $R_{\mu\nu} = 0$.

The curvature equations are themselves a shorthand for ten partial-differential equations. The equations are rich in mathematical interest but are a nightmare to solve. The equations are numerous; worse, they are nonlinear. Therefore, the usual method for handling linear equations – guessing a general form for the solution and making new solutions by combining instances of the general form – does not work. One can spend a decade learning advanced mathematics to solve the equations exactly. Instead, apply a familiar principle: When the going gets tough, lower your standards. By sacrificing some accuracy, we can explain light bending in fewer than one thousand pages – using mathematics and physics that you (and I!) already know.

The simpler method is dimensional analysis, in the usual three steps:

1. Find the relevant parameters.
2. Find dimensionless groups.
3. Use the groups to make the most general dimensionless statement.
4. Add physical knowledge to narrow the possibilities.

These steps are done in the following sections.

5.4.1 Finding parameters

The first step in a dimensional analysis is to decide what physical parameters the bending angle can depend on. For that purpose I often start with an unlabeled diagram, for it prods me into thinking of labels; and many of the labels are parameters of the problem.

This is an interesting approach to solving an abstract problem. It kind of makes sense. It stirs ideas and makes the thinking process more exciting.

I agree. It reminds me of the usefulness of Free Body Diagrams.

It is often hard to visualize all of the parameters of a problem in one's mind, but a diagram can bring many things to light.

I like this because it lets you think about all kinds of variables that could be causing the bend while keeping track of things like direction. Really reminds me of those free body diagrams from 8.01

this is how I started all my physics problems...!

This approach makes a lot of sense, and it is something that I find I do very often since I am a visual learner. It just makes more sense to draw things out.

I think all primates, including humans, are visual learners. We have so much brain hardware devoted to visual processing, so we are very smart when we use pictures and diagrams. People who call themselves visual learners are people who have realized this out already!

labels == quantities?

i think this might refer to variables?

I think it's closer to variables. As in, he is not actually drawing on the force lines and "labeling" which force causes the light to do what. Just a general picture of what it is doing.

He is trying to identify the important dimensions present.

yeah, I read his paper on teaching first year physics in college and how these diagrams are the best to conceptualize the problem, I think that's what he's doing here.

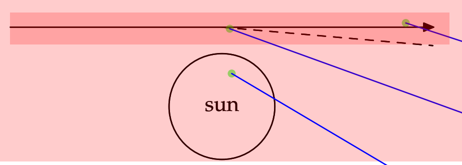

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

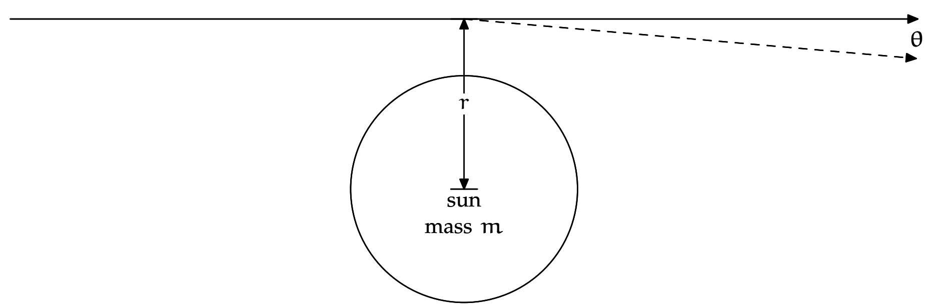

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

Comments on page 2

On a completely unrelated note to this memo, I love the auto adjust when you scroll to a new page.

> Sacha will be very glad to know that! (I'll point him to this thread.)

why is light represented as a straight line? how are you quantifying the amount of light included in this "line"?

What does this line represent?? Is it light? If so where is it coming from if that's the sun? Other stars?

when i read the text leading into the diagram, my assumption was that we were thinking in terms of the earth's gravity...thinking more about it, i get why the sun is _so_ much better, but it might be nice to say sun earlier...or explain why the sun works better.

The paragraph above was kind of lost on me until I saw this diagram - maybe having it before the page break would be good.

we see the light from other stars and galaxies in the night sky. is this same light coming from these distant places assumed to be straight and unaffected by other suns (including our own) or blackholes this far away from their origin?

I would like to see an example of a parameter that you might think it useful but is actually not and the reason for that. I liked in class when you asked us to list parameters in the "Catholic falling into the manure" example. Then, you went through the ones suggested by the class and gave reason as to why they would or would not be useful. This helped me to understand why some parameters are important while others are irrelevant.

Is it bad to include too many parameters? Or will they come out of the dimensional analysis at the end?

> Including too many parameters will make the dimensional analysis very difficult. You'll end up with a statement of the form, "this group is a function of these other five groups", and have a hard time figuring out the function of five variables.
>
> So, toss out parameters early, as if they were deadweight in a lifeboat. If you overdo it, no worries: add one or two back later to fix it up. In short, err on the side of too few than too many.

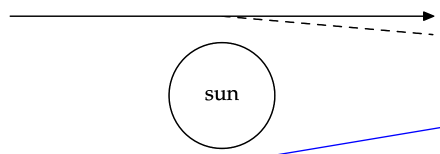

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

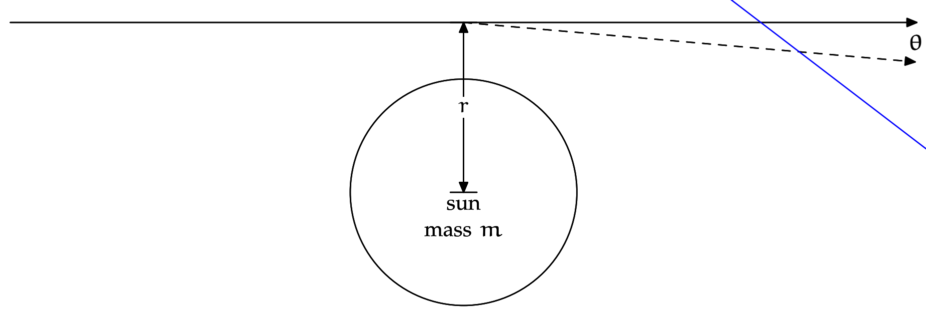

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

I think you meant to add in "are" between "Here" and "various" to read "Here are various..."

Does it help to start with the equations we know about gravity then eliminate ones that aren't relevant? I'm not so good at just thinking about parameters for a situation without an equation or something as a starting point.

I think the point of the dimensional analysis is that we don't really know what the equations are or how to use them. For this situation, we don't know how to calculate the force of gravity on light, what would you use for the mass of light? Dimensional analysis allows us to not worry about the physics.

Right. And if you use the Einstein equations naively, you'll often miss an important quantity, namely G (item 4 in the list), because the equations are often written in a unit system in which c and G are both set to 1.

I think that's pretty much what he's doing here–he's using the diagram to come up with parameters that affect gravity (and these parameters he's listing are in the gravity equations we know). basically he uses the diagram as his starting point.

After making a diagram like this, maybe the light's not having an easy mass is apparent, and so the usual Gm1m2/R^2 doesn't apply. To some degree that's why we're trying this problem, since it baffles me how you would start with Newtonian gravity.

It was nice to see this reasoning worked out. I feel like identifying the relevant parameters is probably the most daunting aspect when using dimensional analysis to solve problems.

We've mentioned this in class, but I oftentimes get too focused on other parameters that I forget the parameter that I want in the end!

that's why it is helpful to list it first always... at least for me.

not a "have to" thing but the sentence could be clearer if it said "Since we are solving for theta, the list obviously must include theta" or something more direct

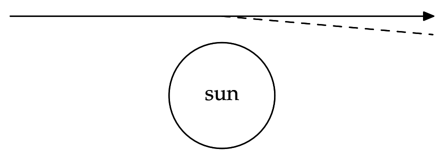

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

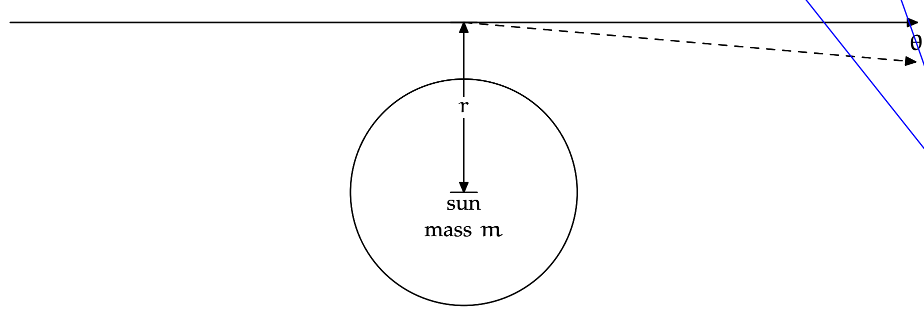

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

how exactly do we simplify the problem to break it down to just theta? doesnt light and all mass want to orbit as opposed to simply bending linearly?

If you think of it in terms of his diagram even without knowing the physics, it's clear that there is a bending angle theta. And uh, still without knowing the physics, but just experience tells us that light will bend without orbiting (it escapes that)

I think it should be stated earlier the relationship between curvature and theta. I, for some reason, thought they were different things, and in the explanation at the start, theta is not mentioned.

Of course this is the classical example, but some motivation for using the sun might be nice.

I concur, at first I read back to see if I missed something on why we chose this particular star of ours instead of some other stellar body.

the sun is the celestial body that is easiest for us to relate to and explain this phenomena. Everyone knows about the sun, while others may not know as much about other celestial bodies.

This question reminds me that there is another reason to use the sun: We move a lot relative to the sun, so we can take pictures of the sky with and without the sun (to do it with the sun requires an eclipse). Whereas relative to other stars, we hardly move at all, so we wouldn't get with-and-without pictures.

If I'm understanding this correctly, perhaps use bullets to represent that they are different groups? Or adding "and"?

I'm not sure what you mean by this.

A numbered list implies a step-by-step process (i.e. the points are ordered) whereas bullets do not.

I can guess that the mass definitely has to do with it, but how large are black holes compared with other objects in our galaxy? I have no intuition about this.

I feel like this sentence doesn't flow too well after the previous one. Perhaps with a better transition or parallel structure, this sentence would flow more smoothly in the text. How about "The huge mass of a black hole is probably the reason it can deflect light."

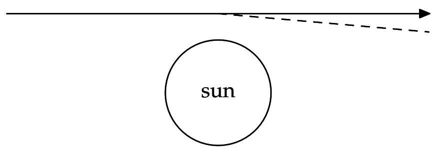

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

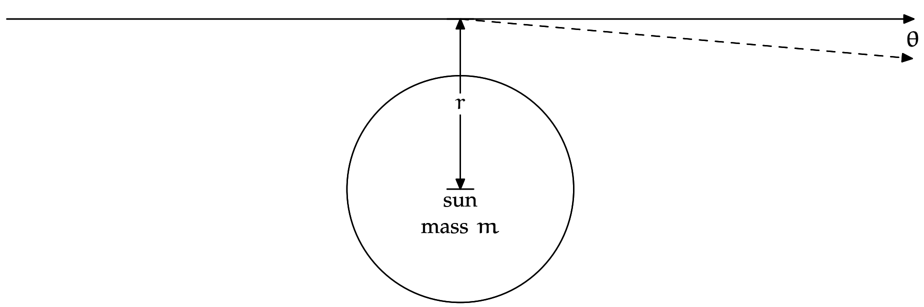

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

Deflect makes sense, but i had to read the sentence twice to understand it. When I think of a black hole, I usually think of it attracting light, not deflecting it, because of its mass. Maybe use the word bend instead?

i agree. i had the same mental image.

I really liked how you included the parameters, so are you saying that light from the sun is bent? Because if the gravitational force of the sun doesn't bend light, why are you including its mass in this calculation?

So we include quantities to solve for and the values that affect those quantities?

whether the mass is concentrated in a star, or if it is concentrated in a point mass of a black hole, it is still the same mass with the same gravitational force.

how do you rationalize using these parameters and not others?

I really like the enumeration and reasoning listed here. It helps me think about these types of problems properly. To answer the above question, I think you just need to limit yourself with a certain number of parameters. Obviously you could go on and on listing other parameters leading to a more precise solution, but that defeats the purpose.

It seems like you need to include things that would affect your first-order solution, and then include things such that the dimensions work out.

Aren't these reasons the rationalization? What others would you suggest?

you could use others if you wanted as long as they make sense. these seem to make logical sense to me, but if you had other ones that appealed to you and made sense I'm confident you would find a similar result.

Now that I think about it, his parameters are not hard to reason. Of course, without the knowledge that gravity bends light, we can assume that light acts like a mass particle and the analogy would still be correct.

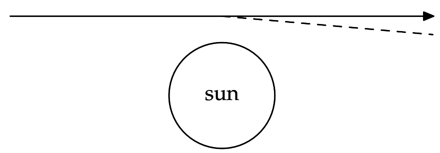

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

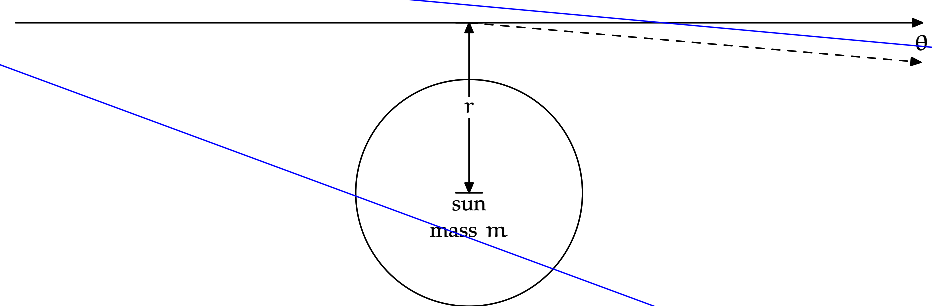

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

technically this still doesn't define r that well, since the light will start bending (slightly) before it reaches what would have been the point of closest approach (i.e. its path isn't quite a line with a sharp corner, as drawn below).

is this the same as the line that's perpendicular to the light ray that goes through the center of the sun?

yes, i think it means that line. The dotted line representing the bending isnt supposed to be straight, but since we are approximating this, we can neglect the curvature.

I wish I had read this earlier. I feel like it would have helped me a lot with other classes.

This personification of dimensional analysis is kind of strange.

Could we be sure though that light uses the same G that affects macroscopic physics?

Can you perhaps explain why you include a constant in your list of parameters? I agree that it's important for dimensional analysis, but it is fundamentally different from the previous 3 items and it seems a little out of place on this list. I would have made a special note or a separate list of related constants in this problem.

I think we have to include G here for the dimensional analysis (otherwise we can't get rid of units if we only look at angle, mass, and distance)

I was also a little confused by including the constant G, although it does seem to make sense I'm not sure where it fits in.

nice diagram

I really appreciate how its a copy of the one just above with the labels added. Something like this would've been useful in the previous sections where we built upon trees and the like but only showed the final diagram.

Yeah, adding pictures in step-by-step is incredibly useful when reading.

Good to know! I'll do that more often. Tufte, a master of information presentation, discusses a similar principle that he calls "Small multiples" (_Envisioning Information_, pp. 28-33).

I really like how this is presented...the same information in 3 different ways! it's awesome.

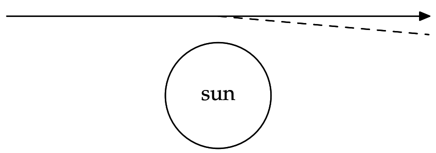

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

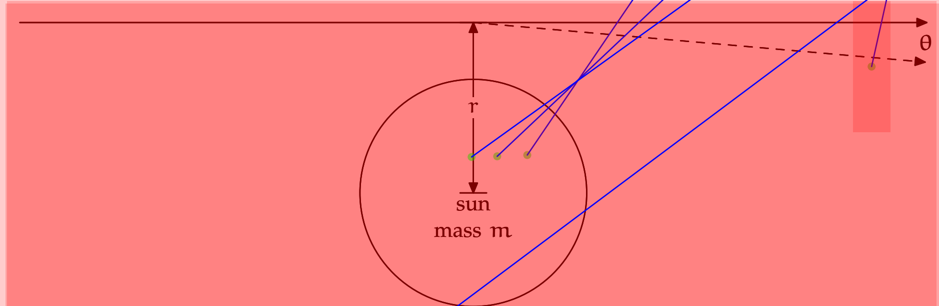

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

I like how this diagram is here to easily see the different parameters.

nice summary

So I've been reading that a photon has been estimated to have a rest mass on the order of 10ˆ-27 eV. Is the gravity acting on that mass? or does it have something to do with the idea that as something's speed approaches the speed of light it's mass becomes infinite?...

Neither, actually. Think of a marble rolling on a rubber sheet: normally it would travel in a straight line, but if a rock is placed on the sheet, it will warp the sheet and cause the marble to change direction as it goes past the rock.

where would you included G, or since it is a constant it is not on the diagram. I just think I would have forgotten it, is there some way to think about these constants and not forget them?

I really like this diagram to show how we are using each parameter.

These types of diagrams are super useful when doing physics problems.

I really like the coupling of the explanatory list, table, and diagram! 3 times is a charm.

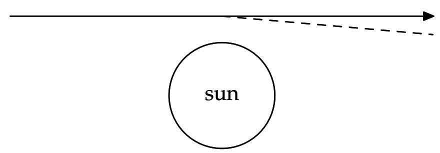

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

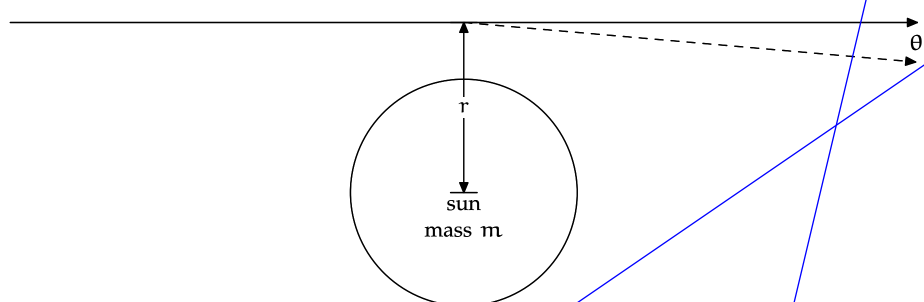

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

is this method less effective when the dimensions are not as simple as those in mechanics?

I've only done problems using mass, length, time, and temperature, but I don't see why you couldn't also use this method for other dimension. What types of things were you thinking of?

why doesn't the angle have a dimension of degrees?

See above, angles are dimensionless. We sort of covered this in class in the inclined plane problem, where gsinTheta was lumped together and had the same units as g.

Here is where the distinction between units and dimensions is helpful. Angles are dimensionless: They are the ratio of two lengths (arc length divided by radius). But they have units: radians, degrees, gradians (my ancient calculator has those). The universe cares about dimensions, not our system of units, so reason using dimensions. Often, reasoning using units and dimensions leads you to the same place; but angles are one area where one must be careful.

the whole reason we "lower our standards" is so that we can use these dimensions that we're familiar with...so i think if you hit a situation with really complicated parameters you'd try to express them in terms of these simple ones...

This was really useful to have tabulated here. Though, I have to ask, is angle really also dimensionless?

I think this goes back to a previous section, where the he discussed a difference between dimensions and units.

Where was that?

This was, I believe, covered in class. The only problem this represents is making me think that since it is dimensionless, I can also arbitrarily multiply it when I try and cancel things.

When I think back on trig, I think of angles as length over length (dimensionless), which is why you can use operations like sine on it. Otherwise you'd have some weird dimension like sine-length.

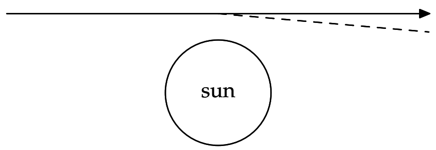

Here various parameters and reasons to include them:

1. The list has to include the quantity to solve for. So the angle θ is the first item in the list.
2. The mass of the sun, m, has to affect the angle. Black holes greatly deflect light, probably because of their huge mass.
3. A faraway sun or black hole cannot strongly affect the path (near the earth light seems to travel straight, in spite of black holes all over the universe); therefore r, the distance from the center of the mass, is a relevant parameter. The phrase 'distance from the center' is ambiguous, since the light is at various distances from the center. Let r be the distance of closest approach.
4. The dimensional analysis needs to know that gravity produces the bending. The parameters listed so far do not create any forces. So include Newton's gravitational constant G.

Here is the diagram with important parameters labeled:

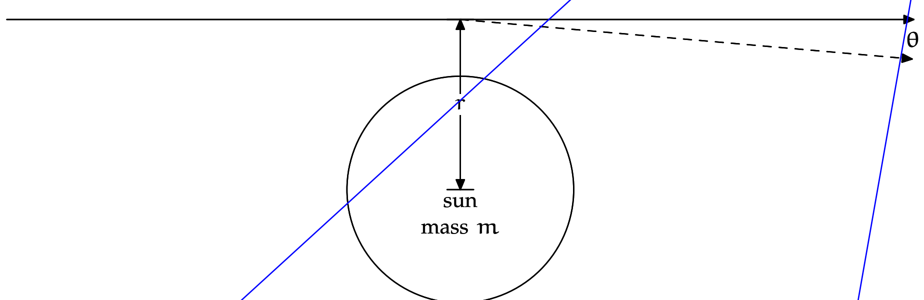

Here is a table of the parameters and their dimensions:

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L

Should not speed be taken into account? would cancel the time values

I thought the same, at this point in the reading. Since it's addressed later, this is answered. However, I wonder why you chose to not include speed from the start. Perhaps, it's just because we're getting too familiar with these sorts of problems from having done a few in previous readings that utilized "c". Maybe a first-time reader of this section wouldn't necessarily think of using "c". Then again, this is the third reading in a 4 part series, people should have read the first 2 and gotten exposure to c already. Why choose a more roundabout way of accounting for c in the third reading?

How do you know that these are all of the parameters you'll need? How do you know when to stop looking for more variables?

I think you can figure it all from the fundamental equations of the system, which in this case, include these parameters.

where L, M, and T represent the dimensions of length, mass, and time, respectively.

5.4.2 Dimensionless groups

The second step is to form dimensionless groups. One group is easy: The parameter θ is an angle, which is already dimensionless. The other variables, G, m, and r, cannot form a second dimensionless group. To see why, following the dimensions of mass. It appears only in G and m, so a dimensionless group would contain the product Gm, which has no mass dimensions in it. But Gm and r cannot get rid of the time dimensions. So there is only one independent dimensionless group, for which θ is the simplest choice.

Without a second dimensionless group, the analysis seems like nonsense. With only one dimensionless group, it must be a constant. In slow motion:

$$\theta = \text{function of other dimensionless groups},$$

but there are no other dimensionless groups, so

$$\theta = \text{constant}.$$

This conclusion is crazy! The angle must depend on at least one of m and r. Let's therefore make another dimensionless group on which θ can depend. Therefore, return to Step 1: Finding parameters. The list lacks a crucial parameter.

What physics has been neglected? Free associating often suggests the missing parameter. Unlike rocks, light is difficult to deflect, otherwise humanity would not have waited until the 1800s to study the deflection, whereas the path of rocks was studied at least as far back as Aristotle and probably for millions of years beforehand. Light travels much faster than rocks, which may explain why light is so difficult to deflect: The gravitational field gets hold of it only for a short time. But none of the parameters distinguish between light and rocks. Therefore let's include the speed of light c. It introduces the fact that we are studying light, and it does so with a useful distinguishing parameter, the speed.

Here is the latest table of parameters and dimensions:

Comments on page 3

Is there any way that this could go on the same page as the chart?

is there any way this info could go in the previous sections? it seems like a lot of this is too late in the text.

this method reminds me a lot of state space equations from 6.011

there seems to be alot of cross over once you get down to such a level

This seems pretty similar to the sliding block on a spring problem we did during lecture...

i'm not sure what this is saying because it doesn't make sense to me grammatically

Should it be "follow the dimensions of mass" ?

Agreed...what does this mean? Even if it was "follow the dimensions of mass," I think it could be made more explicit.

I think he means look at all the terms that have the dimension of mass somewhere in them.

Perhaps this sentence could be reworded to say "To see why, examine those variables which contain mass in their dimensions." Then you could start the next sentence with "Mass appears only in..." It might also be helpful (albeit a bit redundant) to put the dimension of G and m again to show how the product Gm has no mass dimensions in it. But if you clarify the first sentence I mentioned above, it should definitely make this easier to follow.

I don't understand why you want to "follow" mass to begin with...you can see right away that it could be canceled out. why not 'follow' T right from the start?

I feel like this whole part could be worded better for better understanding.

where L, M, and T represent the dimensions of length, mass, and time, respectively.

5.4.2 Dimensionless groups

The second step is to form dimensionless groups. One group is easy: The parameter θ is an angle, which is already dimensionless. The other variables, G, m, and r, cannot form a second dimensionless group. To see why, following the dimensions of mass. It appears only in G and m, so a dimensionless group would contain the product Gm, which has no mass dimensions in it. But Gm and r cannot get rid of the time dimensions. So there is only one independent dimensionless group, for which θ is the simplest choice.

Without a second dimensionless group, the analysis seems like nonsense. With only one dimensionless group, it must be a constant. In slow motion:

$$\theta = \text{function of other dimensionless groups},$$

but there are no other dimensionless groups, so

$$\theta = \text{constant}.$$

This conclusion is crazy! The angle must depend on at least one of m and r. Let's therefore make another dimensionless group on which θ can depend. Therefore, return to Step 1: Finding parameters. The list lacks a crucial parameter.

What physics has been neglected? Free associating often suggests the missing parameter. Unlike rocks, light is difficult to deflect, otherwise humanity would not have waited until the 1800s to study the deflection, whereas the path of rocks was studied at least as far back as Aristotle and probably for millions of years beforehand. Light travels much faster than rocks, which may explain why light is so difficult to deflect: The gravitational field gets hold of it only for a short time. But none of the parameters distinguish between light and rocks. Therefore let's include the speed of light c. It introduces the fact that we are studying light, and it does so with a useful distinguishing parameter, the speed.

Here is the latest table of parameters and dimensions:

what if you accidentally put in speed of light there, then you can form a dimensionless group. so is there a way to check whether this dimensionless group is right or wrong

Agreed. I would have almost certainly included the speed of light. That determines what I would consider the 'duration' that gravity acts on the light. Just like a planet hurtling past will be less affected (smaller theta) at higher speeds, right? (At least intuitively, to me)

Oops, this is addressed entirely at the bottom of this page.

That would be a useful accident, and you'd be well on the way to a solution (see the upcoming text).

this paragraph didnt make any sense to me, but i know what you're trying to say without it.

All it really takes to understand it is to look back at the chart any time a parameter is mentioned. More often than not, you can piece together what's being said based on how the units interact.

don't quite show of this statement? I know you need a second dimensionless group but couldn't i use other methods to finish off the analysis?

Er? I think I'm confused by what this means.

haha I agree...taking it slower so we can understand?

Yeah just a step by step run through of what we're doing.

When I read this I had to smile; I got where you were going immediately, and although I think the earlier explanation is fine, this was a nice summary!

haha yeah, I didn't get it at first, I thought time was going to get slowed down or something lol

While this seems obvious now, esp doing a lot of dimensional analysis in 2.006, I think I've been confused by this before, and why we know this. Perhaps justifying this would help.

I agree, we never really explained why one dimensionless group is a function of the others

where L, M, and T represent the dimensions of length, mass, and time, respectively.

5.4.2 Dimensionless groups

The second step is to form dimensionless groups. One group is easy: The parameter θ is an angle, which is already dimensionless. The other variables, G, m, and r, cannot form a second dimensionless group. To see why, following the dimensions of mass. It appears only in G and m, so a dimensionless group would contain the product Gm, which has no mass dimensions in it. But Gm and r cannot get rid of the time dimensions. So there is only one independent dimensionless group, for which θ is the simplest choice.

Without a second dimensionless group, the analysis seems like nonsense. With only one dimensionless group, it must be a constant. In slow motion:

$$\theta = \text{function of other dimensionless groups},$$

but there are no other dimensionless groups, so

$$\theta = \text{constant}.$$

This conclusion is crazy! The angle must depend on at least one of m and r. Let's therefore make another dimensionless group on which θ can depend. Therefore, return to Step 1: Finding parameters. The list lacks a crucial parameter.

What physics has been neglected? Free associating often suggests the missing parameter. Unlike rocks, light is difficult to deflect, otherwise humanity would not have waited until the 1800s to study the deflection, whereas the path of rocks was studied at least as far back as Aristotle and probably for millions of years beforehand. Light travels much faster than rocks, which may explain why light is so difficult to deflect: The gravitational field gets hold of it only for a short time. But none of the parameters distinguish between light and rocks. Therefore let's include the speed of light c. It introduces the fact that we are studying light, and it does so with a useful distinguishing parameter, the speed.

Here is the latest table of parameters and dimensions:

I think I lost track of the logic here – theta is a function of other dimensionless groups, there are no other dimensionless groups... so why must it be a constant?

I think it means that we were trying to approximate theta as a function of other dimensionless groups in the analysis, but because the only dimensionless group is theta, then there's no other varying "function" that theta could depend on. Therefore, theta must be equal to a constant value.

I'm confused too... Why does theta have to be a function of something without dimensions?

how do you know that you didn't just miss a term? that would make the most sense in this situation.

this part was great! it really showed how to conceptualize the thinking in terms of dimensions well, and what was going on in your head.

!!!!!!!!!!!!!!!!!!! wow. this seems so hard to believe.

so without enough dimensionless groups, it's possible to come to wrong conclusions? How do you know when you have enough dimensionless groups?

yeah i'm confused about this same point too

I think the point to take away from this is that you need at least two dimensionless groups to solve the problem. Some times you may get more than two dimensionless groups, then you should reduce your number of dimensionless groups to a subset which are independent, meaning you can't make any one group in your set from a combination of the other groups.

Without the correct parameters, you will not come to a viable conclusion - you have to identify all of them! If you end up with only 1 and you are sure of it, then you have a very boring problem.

Thanks! This explanation really helped.

Right – in this problem, physical reasoning tells you that you need at least two dimensionless groups. In other problems, one is enough: for example, in estimating how long it takes an object to free fall from a height h (no air resistance).

where L, M, and T represent the dimensions of length, mass, and time, respectively.

5.4.2 Dimensionless groups

The second step is to form dimensionless groups. One group is easy: The parameter θ is an angle, which is already dimensionless. The other variables, G, m, and r, cannot form a second dimensionless group. To see why, following the dimensions of mass. It appears only in G and m, so a dimensionless group would contain the product Gm, which has no mass dimensions in it. But Gm and r cannot get rid of the time dimensions. So there is only one independent dimensionless group, for which θ is the simplest choice.

Without a second dimensionless group, the analysis seems like nonsense. With only one dimensionless group, it must be a constant. In slow motion:

$$\theta = \text{function of other dimensionless groups},$$

but there are no other dimensionless groups, so

$$\theta = \text{constant}.$$

This conclusion is crazy! The angle must depend on at least one of m and r. Let's therefore make another dimensionless group on which θ can depend. Therefore, return to Step 1: Finding parameters. The list lacks a crucial parameter.

What physics has been neglected? Free associating often suggests the missing parameter. Unlike rocks, light is difficult to deflect, otherwise humanity would not have waited until the 1800s to study the deflection, whereas the path of rocks was studied at least as far back as Aristotle and probably for millions of years beforehand. Light travels much faster than rocks, which may explain why light is so difficult to deflect: The gravitational field gets hold of it only for a short time. But none of the parameters distinguish between light and rocks. Therefore let's include the speed of light c. It introduces the fact that we are studying light, and it does so with a useful distinguishing parameter, the speed.

Here is the latest table of parameters and dimensions:

I wouldn't say it's crazy, just very trivial...like concluding 1= 1?

> It's different than 1 = 1, which is indeed trivial. 1=1 is a mathematical tautology, and is always true (just boring). Whereas theta=constant is a physical prediction, and it is almost certainly bogus (although it is interesting).

So i guess the process is iterative. Find some parameters that look like they're relevant. If they don't fit then go back and find more parameters.

This process is still a little ambiguous to me, is there any set of procedures we can use to see if we have found all of the required constants? and what exactly is 'Free Associating'

I agree...free associating may not work for me sometimes. Is there a more formal process to find dimensionless groups?

Ooh yeah I def. agree.

Yeah, I definitely agree. The idea of using the speed of light seems obvious now because you just stated that we should, but I wouldn't have a clue about where to start or how to find other necessary values if I was in this position. Is there a more defined way we can go about finding these parameters?

I think guess and check is really the only way to go about it. The only way to complete the problem formulaicly is through the ugly physics which we're trying to avoid.

Ugly? Really? Maybe extremely difficult, but ugly?

Yeah, this phrase is a bit ambiguous as to the meaning of this sentence.

Could we go over free associating in class?

Random and awesome.

hahaha, true and a cute observation

where L, M, and T represent the dimensions of length, mass, and time, respectively.

5.4.2 Dimensionless groups

The second step is to form dimensionless groups. One group is easy: The parameter θ is an angle, which is already dimensionless. The other variables, G, m, and r, cannot form a second dimensionless group. To see why, following the dimensions of mass. It appears only in G and m, so a dimensionless group would contain the product Gm, which has no mass dimensions in it. But Gm and r cannot get rid of the time dimensions. So there is only one independent dimensionless group, for which θ is the simplest choice.

Without a second dimensionless group, the analysis seems like nonsense. With only one dimensionless group, it must be a constant. In slow motion:

$$\theta = \text{function of other dimensionless groups,}$$

but there are no other dimensionless groups, so

$$\theta = \text{constant.}$$

This conclusion is crazy! The angle must depend on at least one of m and r. Let's therefore make another dimensionless group on which θ can depend. Therefore, return to Step 1: Finding parameters. The list lacks a crucial parameter.

What physics has been neglected? Free associating often suggests the missing parameter. Unlike rocks, light is difficult to deflect, otherwise humanity would not have waited until the 1800s to study the deflection, whereas the path of rocks was studied at least as far back as Aristotle and probably for millions of years beforehand. Light travels much faster than rocks, which may explain why light is so difficult to deflect: The gravitational field gets hold of it only for a short time. But none of the parameters distinguish between light and rocks. Therefore let's include the speed of light c. It introduces the fact that we are studying light, and it does so with a useful distinguishing parameter, the speed.

Here is the latest table of parameters and dimensions:

I don't think humans existed millions of years ago... So they probably weren't studying rocks.

> However, our ancestors were studying rocks, at least implicitly. The primates who had bad abilities at predicting didn't do so well and are less likely to be our ancestors. So, evolution studied it (if I can personify evolution), and built some of that knowledge into our brains.

I'd move this sentence between "... rock." and "Therefore ..." ... it makes more sense/flows better there.

Never thought of it like that!

This is cool - turning light into a projectile

Yeah I never would have thought about light in this way, definitely a cool way to look at it.

this goes back to the constant thing I just talked about-for some reason I feel like they shouldn't matter since they are constant for all while the other variables are the one that are changing and would actually have an effect on the output

now it seems so obvious why did we not include this.

it seems obvious and this shows one of the shortcomings of our diagram - it showed light as a line and not some moving thing so we didnt think to add c to the list of parameters

A little confusing on why c is added in the section of dimensionless constants? Shouldn't it be somewhere else?

I think the addition of c brings the size of the differences into perspective. Also, it works perfectly because we are analyzing the speed of light.

where L, M, and T represent the dimensions of length, mass, and time, respectively.

5.4.2 Dimensionless groups

The second step is to form dimensionless groups. One group is easy: The parameter θ is an angle, which is already dimensionless. The other variables, G, m, and r, cannot form a second dimensionless group. To see why, following the dimensions of mass. It appears only in G and m, so a dimensionless group would contain the product Gm, which has no mass dimensions in it. But Gm and r cannot get rid of the time dimensions. So there is only one independent dimensionless group, for which θ is the simplest choice.

Without a second dimensionless group, the analysis seems like nonsense. With only one dimensionless group, it must be a constant. In slow motion:

$$\theta = \text{function of other dimensionless groups},$$

but there are no other dimensionless groups, so

$$\theta = \text{constant}.$$

This conclusion is crazy! The angle must depend on at least one of m and r. Let's therefore make another dimensionless group on which θ can depend. Therefore, return to Step 1: Finding parameters. The list lacks a crucial parameter.

What physics has been neglected? Free associating often suggests the missing parameter. Unlike rocks, light is difficult to deflect, otherwise humanity would not have waited until the 1800s to study the deflection, whereas the path of rocks was studied at least as far back as Aristotle and probably for millions of years beforehand. Light travels much faster than rocks, which may explain why light is so difficult to deflect: The gravitational field gets hold of it only for a short time. But none of the parameters distinguish between light and rocks. Therefore let's include the speed of light c. It introduces the fact that we are studying light, and it does so with a useful distinguishing parameter, the speed.

Here is the latest table of parameters and dimensions:

I feel dumb for not guessing that

Yeah, i was wondering in the above section on finding parameters, whether or not we needed to include the speed of light. but then I thought to myself that the speed of light is constant for all light, so it shouldn't change anything. Except in this case we need to include it to find a dimensionless group.

yeah from this point, it seems that some variable with speed should have been included, I think it's great that you left this to be concluded from the second group's lack of dimensionlessness.

So do we need to begin considering light as mass?

yeah, I was thinking, why do we chose c (the speed of light) to further characterize the phenomena rather than using the mass of light?

edit: updated

Parameter	Meaning	Dimensions
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L
c	speed of light	LT^{-1}

Length is strewn all over the parameters (it's in G, r, and c). Mass, however, appears in only G and m, so the combination Gm cancels out mass. Time also appears in only two parameters: G and c. To cancel out time, form Gm/c^2. This combination contains one length, so a dimensionless group is Gm/rc^2.

5.4.3 Drawing conclusions

The most general relation between the two dimensionless groups is

$$\theta = f\left(\frac{Gm}{rc^2}\right).$$

Dimensional analysis cannot determine the function f, but it has told us that f is a function only of Gm/rc^2 and not of the variables separately.

Physical reasoning and symmetry narrow the possibilities. First, strong gravity – from a large G or m – should increase the angle. So f should be an increasing function. Now apply symmetry. Imagine a world where gravity is repulsive or, equivalently, where the gravitational constant is negative. Then the bending angle should be negative; to make that happen, f must be an odd function: namely, $f(-x) = -f(x)$. This symmetry argument eliminates choices like $f(x) \sim x^2$.

The simplest guess is that f is the identity function: $f(x) \sim x$. Then the bending angle is

$$\theta = \frac{Gm}{rc^2}.$$

But there is probably a dimensionless constant in f. For example,

$$\theta = 7\frac{Gm}{rc^2}$$

and

Comments on page 4

Why is t his not degrees or radians?

Degrees and radians are still dimensionless

The table is very helpful. (allows us to keep track of the dimensionless parameters).

It seems like these constants were constructed for the sole purpose of generating dimensionless groups. Like people were trying to think of how quantities were related, and just thought of a new quantity that would conveniently cancel out any bothersome dimensions. How, for example, were the dimensions of G found?

> You find the dimensions of G from any other equation that you trust with G in it, e.g. F=GMm/r^2.

Why don't we consider the mass of light? Doesn't the gravitation equation require two masses? Would account for this mass, even though it is incredibly small, change the proportionality of the equation?

I think we ignore it because the mass cancels out in Gm anyway.

Also, that may be getting into the nitty-gritty physics which we are trying to avoid.

Correct, in lowering on standards on accuracy, we can ignore the really picky and more or less insignificant quantities that would be just extra work.

so are we omitting it because the mass is insignificantly small or because it cancels out anyways?

I would say that the mass of light is relatively insignificant and like said above by lowering accuracy we can leave out some things. It doesn't look like the mass actually cancels out though since here we've just made a non-dimensional equation, not canceled anything out.

> In gravity problems, the mass always cancels out (the force is GMm/r^2 and the acceleration is GM/r^2, which is independent of m). I think I should discuss this point earlier in the text when the list is being constructed.

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L
c	speed of light	LT^{-1}

Length is strewn all over the parameters (it's in G, r, and c). Mass, however, appears in only G and m, so the combination Gm cancels out mass. Time also appears in only two parameters: G and c. To cancel out time, form Gm/c^2. This combination contains one length, so a dimensionless group is Gm/rc^2.

5.4.3 Drawing conclusions

The most general relation between the two dimensionless groups is

$$\theta = f\left(\frac{Gm}{rc^2}\right).$$

Dimensional analysis cannot determine the function f, but it has told us that f is a function only of Gm/rc^2 and not of the variables separately.

Physical reasoning and symmetry narrow the possibilities. First, strong gravity – from a large G or m – should increase the angle. So f should be an increasing function. Now apply symmetry. Imagine a world where gravity is repulsive or, equivalently, where the gravitational constant is negative. Then the bending angle should be negative; to make that happen, f must be an odd function: namely, $f(-x) = -f(x)$. This symmetry argument eliminates choices like $f(x) \sim x^2$.

The simplest guess is that f is the identity function: $f(x) \sim x$. Then the bending angle is

$$\theta = \frac{Gm}{rc^2}.$$

But there is probably a dimensionless constant in f. For example,

$$\theta = 7\frac{Gm}{rc^2}$$

and

I really like this paragraph. I like how it goes through the process step by step. It cleared up my questions when it came to dimensionless groups.

Couldn't you have written the inverse of rc^2/Gm?

include a table of the dimensionless constants since there are 2 of them?

The previous two sections are very good; the explanation is clear and following is easy. It feels very similar to lecture, which is also easy to follow. I wouldn't alter it too much.

This is so cool how you can figure out how to put together the equation. But I still don't understand if we are treating light as a mass or not?

I was really surprised you could come to this conclusion without any calculus or physics.

I took me reading this a few times to see the 'function of' notation there...kept thinking, why aren't you saying this and then proving it.

i'm confused by this function business.

Well, theta could equal Gm/rc^2, or the sin(theta) could, or we could square the entire right side, or...

All we know is that the angle relates to some function of that "group" although we do not know if it is linear with the group, etc.

I really liked the example in class where you showed us a complex physics phenomena where you don't need to memorize individual quantities but only their combination. It's really cool to see the parallels between this text and class and helps to solidify my understanding of these concepts.

do you have any idea of what the function is? is it usually linear, or polynomial or could it really be anything? is that something you become better with over time?

I'm a bit confused, if we know from earlier that the end result is non-linear, do these analysis still hold?

The non-linearity may come in in the form of the function f, but we can still say that it must be a function of this particular product, I think.

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L
c	speed of light	LT^{-1}

Length is strewn all over the parameters (it's in G, r, and c). Mass, however, appears in only G and m, so the combination Gm cancels out mass. Time also appears in only two parameters: G and c. To cancel out time, form Gm/c^2. This combination contains one length, so a dimensionless group is Gm/rc^2.

5.4.3 Drawing conclusions

The most general relation between the two dimensionless groups is

$$\theta = f\left(\frac{Gm}{rc^2}\right).$$

Dimensional analysis cannot determine the function f, but it has told us that f is a function only of Gm/rc^2 and not of the variables separately.

Physical reasoning and symmetry narrow the possibilities. First, strong gravity – from a large G or m – should increase the angle. So f should be an increasing function. Now apply symmetry. Imagine a world where gravity is repulsive or, equivalently, where the gravitational constant is negative. Then the bending angle should be negative; to make that happen, f must be an odd function: namely, $f(-x) = -f(x)$. This symmetry argument eliminates choices like $f(x) \sim x^2$.

The simplest guess is that f is the identity function: $f(x) \sim x$. Then the bending angle is

$$\theta = \frac{Gm}{rc^2}.$$

But there is probably a dimensionless constant in f. For example,

$$\theta = 7\frac{Gm}{rc^2}$$

and

this answers my previous question

I wasn't sure how we would go about determining the functions of out dimensionless values but this makes a lot of sense. I was confused about this in class a bit so I'm glad we're covering it now.

meaning it has a positive slope? but what parameters would be the independent one?

I would have stopped here, but like when you make me realize how much further I can take my estimation.

In this case is negative theta a "pushing away" effect, like the force would cause the angle of deflection to be away from the sun?

but it shouldn't necessarily be equal. Since the bending occurs over a window of time, the object bending towards the sun will be closer (and therefore more affected) by gravity at the end of that timeframe than the object bending away.

The extreme example, is that if the object is moving too slowly, it will spiral into the sun. There is no equivalent 'spiraling away', only being repelled somewhat linearly.

My point is that f(-x) shouldn't intuitively equal f(x), at least not for large-ish x.

I'm not sure I understand what it means for the bending angle to be negative? Also, the point given by the student above is also quite interesting.

But what if you think of the sun as being on the opposite side of the light ray? This would demonstrate a negative gravity from the perspective below the light ray, and also show a negative theta, where the light would bend up, away from anything below the ray.

I like this explanation of why it can't be an odd function.

no...it can't be an even function. maybe the explanation didn't work that well...

Parameter	*Meaning*	*Dimensions*
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L
c	speed of light	LT^{-1}

Length is strewn all over the parameters (it's in G, r, and c). Mass, however, appears in only G and m, so the combination Gm cancels out mass. Time also appears in only two parameters: G and c. To cancel out time, form Gm/c^2. This combination contains one length, so a dimensionless group is Gm/rc^2.

5.4.3 Drawing conclusions

The most general relation between the two dimensionless groups is

$$\theta = f\left(\frac{Gm}{rc^2}\right).$$

Dimensional analysis cannot determine the function f, but it has told us that f is a function only of Gm/rc^2 and not of the variables separately.

Physical reasoning and symmetry narrow the possibilities. First, strong gravity – from a large G or m – should increase the angle. So f should be an increasing function. Now apply symmetry. Imagine a world where gravity is repulsive or, equivalently, where the gravitational constant is negative. Then the bending angle should be negative; to make that happen, f must be an odd function: namely, $f(-x) = -f(x)$. This symmetry argument eliminates choices like $f(x) \sim x^2$.

The simplest guess is that f is the identity function: $f(x) \sim x$. Then the bending angle is

$$\theta = \frac{Gm}{rc^2}.$$

But there is probably a dimensionless constant in f. For example,

$$\theta = 7\frac{Gm}{rc^2}$$

and

I like this reasoning to determine f(x)

Yeah, this is really interesting reasoning. I never would have thought of this, but it makes total sense.

Can we go over an example where f(x) /= x? It seems like for everything we have done so far we can approximate them as the same and I would be interested to see and example where physical reasoning leads us to another answer

Yeah I really enjoy this. I would have never thought of using this sort of thinking. It's extremely useful.

It is a good guess, but I think it'd be nice to add why guessing x^3 doesn't make sense. I know it's pretty intuitive, but a sentence or two would make the thinking clear.

In this case, I think the example actually detracts. My first instinct was that the correct answer had a 7 in it.

It might be better to write "const." x Gm/(rc^2) like you do on the board sometimes.

Disagree. I would have no idea what constant to use so 7 or 0.3 seem equally likely. Although the confusion about the twiddle versus the proportional sign below does confuse me; I thought twiddle also meant "approximately" and not "proportional to"

Proportional to means the dimensions don't have to agree. A single twiddle means that the dimensions agree but there is a constant of proportionality. A double twiddle means approximately equal (the constant of proportionality is about 1)

I think I shouldn't use 7. It's such a friendly number; not only is it, in popular culture, a lucky number, it is also not irrational. So, I'll use all implausible constants.

Parameter	Meaning	Dimensions
θ	angle	–
m	mass of sun	M
G	Newton's constant	$L^3T^{-2}M^{-1}$
r	distance from center of sun	L
c	speed of light	LT^{-1}

Length is strewn all over the parameters (it's in G, r, and c). Mass, however, appears in only G and m, so the combination Gm cancels out mass. Time also appears in only two parameters: G and c. To cancel out time, form Gm/c^2. This combination contains one length, so a dimensionless group is Gm/rc^2.

5.4.3 Drawing conclusions

The most general relation between the two dimensionless groups is

$$\theta = f\left(\frac{Gm}{rc^2}\right).$$

Dimensional analysis cannot determine the function f, but it has told us that f is a function only of Gm/rc^2 and not of the variables separately.

Physical reasoning and symmetry narrow the possibilities. First, strong gravity – from a large G or m – should increase the angle. So f should be an increasing function. Now apply symmetry. Imagine a world where gravity is repulsive or, equivalently, where the gravitational constant is negative. Then the bending angle should be negative; to make that happen, f must be an odd function: namely, $f(-x) = -f(x)$. This symmetry argument eliminates choices like $f(x) \sim x^2$.

The simplest guess is that f is the identity function: $f(x) \sim x$. Then the bending angle is

$$\theta = \frac{Gm}{rc^2}.$$

But there is probably a dimensionless constant in f. For example,

$$\theta = 7\frac{Gm}{rc^2}$$

and

maybe instead of having an actual number it might be better to have a "k"?

I don't think it matters what number/variable he uses, since it's an arbitrary constant. It will be removed anyways with the use of twiddle later.

I think that gives you more of a reason to use "k" here rather than an actual number. Or maybe, after using k = 7 and then k = 0.3, you could put a third equation using just k. Then it'll flow nicely into the twiddle later.

I was going to nitpick on the 7 as well. Being an MIT student, I prefer constants in letter form rather than specific cases. But then I read a little farther and saw that he addressed it, so I think in this case it's fine... it appeals both to people who like specificity and generality.

Should this be 'or'?

these could be put on the same line, just saying

$$\theta = 0.3\frac{Gm}{rc^2}$$

are also possible. This freedom means that we should use a twiddle rather than an equals sign:

$$\theta \sim \frac{Gm}{rc^2}.$$

5.4.4 Comparison with exact calculations

All theories of gravity have the same form for the result, namely

$$\theta \sim \frac{Gm}{rc^2}.$$

The difference among the theories is in the value for the missing dimensionless constant:

$$\theta = \frac{Gm}{rc^2} \times \begin{cases} 1 & \text{(simplest guess);} \\ 2 & \text{(Newtonian gravity);} \\ 4 & \text{(Einstein's theory).} \end{cases}$$

Here is a rough explanation of the origin of those constants. The 1 for the simplest guess is just the simplest possible guess. The 2 for Newtonian gravity is from integrating angular factors like cosine and sine that determine the position of the photon as it moves toward and past the sun.

The most interesting constant is the 4 for general relativity, which is double the Newtonian value. The fundamental reason for the factor of 2 is that special relativity puts space and time on an equal footing to make spacetime. The theory of general relativity builds on special relativity by formulating gravity as curvature of spacetime. Newton's theory is the limit of general relativity that considers only time curvature; but general relativity also handles the space curvature. Most objects move much slower than the speed of light, so they move much farther in time than in space and see mostly the time curvature. For those objects, the Newtonian analysis is fine. But light moves at the speed of light, and it therefore sees equal amounts of space and time curvature; so its trajectory bends twice as much as the Newtonian theory predicts.

5.4.5 Numbers!

At the surface of the Earth, the dimensionless gravitational strength is

does this mean that it could be multiplied by any dimensionless constant

I sincerely hope that's the real word to describe that shape. I don't remember your using it in class the day you went over the four symbols, and it seems too fun to be real, but I could have just missed that part.

It's technically called a tilde, but I think twiddle is an accepted 'nickname'. However, i think tilde is more universally known...and maybe it should be changed to that.

And why not a proportional sign, again? This seems to be exactly the time to use a proportional sign – when we are explicitly omitting a constant of proportionality. I feel like there is a lack of consistency throughout these readings on this point.

I thought the twiddle meant approximately, which is roughly the same as the proportional sign. The proportional sign is confusing because it looks like an alpha.

A proportional sign hides a dimensioned constant. We have none missing here. A twiddle hides a dimensionless constant.

This is mentioned a few sentences later

I had not realized the reason for this distinction before. I've been treating twiddles as equivalent to proportionality signs up until now. It would've been nice to explain the difference much earlier on (unless I didn't see it when it was mentioned).

we've been approximating all along. why bring this up now?

I agree, I think we talked about this point in class at some time but it would have been nice to see earlier in writing.

but you don't really know whether the exponent is really 1 though, it might be 2, 3,... how can you assume it's 1 and use a twiddle here

Well it can't be 2... since that's an even function. But I agree otherwise.

That's a very good point that didn't cross my mind. How can we tell if it's 1, 3, etc?

He said earlier that he's just using x^1 because it's the most simple even function. It could in theory be 3,5,etc. but that complexity would take away from the approximation I think.

$$\theta = 0.3\frac{Gm}{rc^2}$$

are also possible. This freedom means that we should use a twiddle rather than an equals sign:

$$\theta \sim \frac{Gm}{rc^2}.$$

5.4.4 Comparison with exact calculations

All theories of gravity have the same form for the result, namely

$$\theta \sim \frac{Gm}{rc^2}.$$

The difference among the theories is in the value for the missing dimensionless constant:

$$\theta = \frac{Gm}{rc^2} \times \begin{cases} 1 & \text{(simplest guess);} \\ 2 & \text{(Newtonian gravity);} \\ 4 & \text{(Einstein's theory).} \end{cases}$$

Here is a rough explanation of the origin of those constants. The 1 for the simplest guess is just the simplest possible guess. The 2 for Newtonian gravity is from integrating angular factors like cosine and sine that determine the position of the photon as it moves toward and past the sun.

The most interesting constant is the 4 for general relativity, which is double the Newtonian value. The fundamental reason for the factor of 2 is that special relativity puts space and time on an equal footing to make spacetime. The theory of general relativity builds on special relativity by formulating gravity as curvature of spacetime. Newton's theory is the limit of general relativity that considers only time curvature; but general relativity also handles the space curvature. Most objects move much slower than the speed of light, so they move much farther in time than in space and see mostly the time curvature. For those objects, the Newtonian analysis is fine. But light moves at the speed of light, and it therefore sees equal amounts of space and time curvature; so its trajectory bends twice as much as the Newtonian theory predicts.

5.4.5 Numbers!

At the surface of the Earth, the dimensionless gravitational strength is

What do you mean "all theories?" Are you just referring to the various ways we describe gravity?

I think he's referring to the fact that all the ways that we describe gravity are just theories because we don't really have any way of knowing if these theories are correct. We just come up with theories and their corresponding formulas to describe what we observe, but we can't prove which one if correct. Therefore, they're all theories.

Sorry – I meant, "all of the proposed theories of gravity", including Newtonian gravity, general relativity, and the Brans-Dicke theory. Actually, I probably should not say "all", since I don't know all of them, and there may be some pretty crazy ones out there. But the list above is a reasonable subset of the reasonable theories.

Is this really the answer we would get after solving the original equation and the subsequent 10 non linear equations? That's crazy that we could arrive at the same thing after going through a page of dimensional analysis

The factor of 4 is what you get after solving the horrible equations. So, the dimensional analysis with a bit of guessing (for the function f) gives you almost all of the GR solution, just not the actual value of the constant.

by simplest guess do you mean what would make it easiest? or did someone say that it was 1

Where did these numbers come from?

I was confused too before reading the paragraph below.

They are constants from different theories. It's a favorite story in physics classes about how Newton was off by a factor of two.

At least to me, this is a silly sentence construction.

Agreed - why is this sentence even here?

I think the sentence is here because he's listing out the reasons for each of the constants (simplest guess, newtonian, ...) and it would seem strange to leave one of them out. But I agree it could use rewording. Even placing "simplest guess" in quotes to show that we're referring to the named constant in the expression above could be helpful.

$$\theta = 0.3\frac{Gm}{rc^2}$$

are also possible. This freedom means that we should use a twiddle rather than an equals sign:

$$\theta \sim \frac{Gm}{rc^2}.$$

5.4.4 Comparison with exact calculations

All theories of gravity have the same form for the result, namely

$$\theta \sim \frac{Gm}{rc^2}.$$

The difference among the theories is in the value for the missing dimensionless constant:

$$\theta = \frac{Gm}{rc^2} \times \begin{cases} 1 & \text{(simplest guess);} \\ 2 & \text{(Newtonian gravity);} \\ 4 & \text{(Einstein's theory).} \end{cases}$$

Here is a rough explanation of the origin of those constants. The 1 for the simplest guess is just the simplest possible guess. The 2 for Newtonian gravity is from integrating angular factors like cosine and sine that determine the position of the photon as it moves toward and past the sun.

The most interesting constant is the 4 for general relativity, which is double the Newtonian value. The fundamental reason for the factor of 2 is that special relativity puts space and time on an equal footing to make spacetime. The theory of general relativity builds on special relativity by formulating gravity as curvature of spacetime. Newton's theory is the limit of general relativity that considers only time curvature; but general relativity also handles the space curvature. Most objects move much slower than the speed of light, so they move much farther in time than in space and see mostly the time curvature. For those objects, the Newtonian analysis is fine. But light moves at the speed of light, and it therefore sees equal amounts of space and time curvature; so its trajectory bends twice as much as the Newtonian theory predicts.

5.4.5 Numbers!

At the surface of the Earth, the dimensionless gravitational strength is

Does the Newtonian approach work here? We can't use the mass... I guess we are really just deflecting the momentum of the photon?

Yeah, that's my interpretation as well. Kind of interesting to see the differences in the factors.

I definitely did not know this- that's REALLY cool

Yeah it is. And although it is something that is extremely complex, this paragraph does a good job of explaining it to non physics majors like myself.

cool!

A diagram illustrating how spacetime is curved might be helping here. I'm thinking in particular of the classic "bowling ball on a sheet" illustration that represents the warping of spacetime around a sun/planet.

it's almost comical to try to explain special relativity in an approximation textbook in 7 sentences.

But it is interesting and relevant - we don't need a perfect understanding but a general background is helpful.

Regardless of how funny you might think it is, I think it's really nice that Sanjoy puts the background of different fields into his book.

I am simply amazed at how much he knows...

How are units compared here? I feel like we just had a section talking about how we can't make this statement.

I'm also confused, what does it mean to move "farther" in time than in space?

Does everything move the same in time? If so that would explain why slower objects move less far in space than in time.

this is really fascinating.

I think this could be reworded to sound less redundant.

$$\theta = 0.3\frac{Gm}{rc^2}$$

are also possible. This freedom means that we should use a twiddle rather than an equals sign:

$$\theta \sim \frac{Gm}{rc^2}.$$

5.4.4 Comparison with exact calculations

All theories of gravity have the same form for the result, namely

$$\theta \sim \frac{Gm}{rc^2}.$$

The difference among the theories is in the value for the missing dimensionless constant:

$$\theta = \frac{Gm}{rc^2} \times \begin{cases} 1 & \text{(simplest guess);} \\ 2 & \text{(Newtonian gravity);} \\ 4 & \text{(Einstein's theory).} \end{cases}$$

Here is a rough explanation of the origin of those constants. The 1 for the simplest guess is just the simplest possible guess. The 2 for Newtonian gravity is from integrating angular factors like cosine and sine that determine the position of the photon as it moves toward and past the sun.

The most interesting constant is the 4 for general relativity, which is double the Newtonian value. The fundamental reason for the factor of 2 is that special relativity puts space and time on an equal footing to make spacetime. The theory of general relativity builds on special relativity by formulating gravity as curvature of spacetime. Newton's theory is the limit of general relativity that considers only time curvature; but general relativity also handles the space curvature. Most objects move much slower than the speed of light, so they move much farther in time than in space and see mostly the time curvature. For those objects, the Newtonian analysis is fine. But light moves at the speed of light, and it therefore sees equal amounts of space and time curvature; so its trajectory bends twice as much as the Newtonian theory predicts.

5.4.5 Numbers!

At the surface of the Earth, the dimensionless gravitational strength is

This is pretty cool.

I would never have guessed that with my naive understanding of physics.

Huh, I didn't realize it worked that way (I still don't understand the physics entirely, but I didn't realize it just doubles like that. I'm sure there's a lot of hidden math here though...)

SO which one is right?!

I don't know much about special relativity or advanced physics but this was a very good and clear explanation. However, how do we know that this is actually the reasoning. Maybe I'm just a skeptic but it would be nice to have some references here.

$$\frac{Gm}{rc^2} \sim \frac{6.7 \cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 6.0 \cdot 10^{24}\,\mathrm{kg}}{6.4 \cdot 10^{6}\,\mathrm{m} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}}} \sim 10^{-9}.$$

This miniscule value is the bending angle (in radians). If physicists want to show that light bends, they had better look beyond the earth! That statement is based on another piece of dimensional analysis and physical reasoning, whose result is quoted without proof: A telescope with mirror of diameter d can resolve angles roughly as small as λ/d, where λ is the wavelength of light; this result is based on the same physics as the diffraction pattern on a CDROM (Section 1.1). One way to measure the bending of light is to measure the change in position of the stars. A lens that could resolve an angle of 10^{-9} has a diameter of at least

$$d \sim \lambda/\theta \sim \frac{0.5 \cdot 10^{-6}\,\mathrm{m}}{10^{-9}} \sim 500\,\mathrm{m}.$$

Large lenses warp and crack; one of the largest existing lenses has $d \sim 6\,\mathrm{m}$. No practical mirror can have $d \sim 500\,\mathrm{m}$, and there is no chance of detecting a deflection angle of 10^{-9}.

Physicists therefore searched for another source of light bending. In the solar system, the largest mass is the sun. At the surface of the sun, the field strength is

$$\frac{Gm}{rc^2} \sim \frac{6.7 \cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 2.0 \cdot 10^{30}\,\mathrm{kg}}{7.0 \cdot 10^{8}\,\mathrm{m} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}}} \sim 2.1 \cdot 10^{-6} \approx 0.4''.$$

This angle, though small, is possible to detect: The required lens diameter is roughly

$$d \sim \lambda/\theta \sim \frac{0.5 \cdot 10^{-6}\,\mathrm{m}}{2.1 \cdot 10^{-6}} \sim 20\,\mathrm{cm}.$$

The eclipse expedition of 1919, led by Arthur Eddington of Cambridge University, tried to measure exactly this effect.

For many years – between 1909 and 1916 – Einstein believed that a correct theory of gravity would predict the Newtonian value, which turns out to be 0.87 arcseconds for light that grazes the surface of the sun. The German mathematician Soldner derived the same result in 1803. Fortunately for Einstein's reputation, the eclipse expeditions that went to test his (and Soldner's) prediction got rained or clouded out. By the time an expedition got lucky with the weather (Eddington's in 1919), Einstein had invented

Comments on page 6

this is really small, but then again what if theta = (Gm/rc^2) ^(- power, then the angle would be big how can you make sure the exponent is not negative i am a bit confused about this

That wouldn't make sense. We know it has to be a positive power because of proportionality. If G increased, then theta would also have to increase.

since you're looking at earth first, shouldn't your diagram on pg97 say earth? ...this would work for my prev. comment too.

I'd like to see a "theta " at the beginning here.

While the rest of this section seems sort of like a long side note about our approximation, it was very interesting! I'm glad I read it and learned a few things about the history of gravity and general relativity.

this should probably be explained at the beginning (if at all) since we started out saying it's a dimensionless value. otherwise, we know it's radians.

actually, this makes me wonder, how did we decide it was in radians and not degrees?

doesn't light bending also have a lot to do with density of the atmosphere and such? are we just disregarding that?

If your question is about whether telescopes on earth will see light bending more due to the atmosphere than due to gravity, then the answer has to be yes, because no telescopes that exist can see bending of light due to gravity anyway.

If your question is a general question about whether we're ignoring atmosphere throughout the whole section, the answer is still yes. The atmosphere is relatively thin (thin as in height above the surface of the earth, not density) and most of the light that passes by the planet won't go through the atmosphere where it's dense enough to matter. And the light that does go through the atmosphere gets diffracted so much that it ceases to matter anyway. More importantly, that's just not the effect we're interested in studying here.

I don't really see how measuring this would work, what is the change being compared to?

I'm a little confused on how they actually measure this, like how does measuring the change in position of the stars identify the angle?

$$\frac{Gm}{rc^2} \sim \frac{6.7 \cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 6.0 \cdot 10^{24}\,\mathrm{kg}}{6.4 \cdot 10^6\,\mathrm{m} \times 3.0 \cdot 10^8\,\mathrm{m\,s^{-1}} \times 3.0 \cdot 10^8\,\mathrm{m\,s^{-1}}} \sim 10^{-9}.$$

This miniscule value is the bending angle (in radians). If physicists want to show that light bends, they had better look beyond the earth! That statement is based on another piece of dimensional analysis and physical reasoning, whose result is quoted without proof: A telescope with mirror of diameter d can resolve angles roughly as small as λ/d, where λ is the wavelength of light; this result is based on the same physics as the diffraction pattern on a CDROM (Section 1.1). One way to measure the bending of light is to measure the change in position of the stars. A lens that could resolve an angle of 10^{-9} has a diameter of at least

$$d \sim \lambda/\theta \sim \frac{0.5 \cdot 10^{-6}\,\mathrm{m}}{10^{-9}} \sim 500\,\mathrm{m}.$$

Large lenses warp and crack; one of the largest existing lenses has $d \sim 6\,\mathrm{m}$. No practical mirror can have $d \sim 500\,\mathrm{m}$, and there is no chance of detecting a deflection angle of 10^{-9}.

Physicists therefore searched for another source of light bending. In the solar system, the largest mass is the sun. At the surface of the sun, the field strength is

$$\frac{Gm}{rc^2} \sim \frac{6.7 \cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 2.0 \cdot 10^{30}\,\mathrm{kg}}{7.0 \cdot 10^8\,\mathrm{m} \times 3.0 \cdot 10^8\,\mathrm{m\,s^{-1}} \times 3.0 \cdot 10^8\,\mathrm{m\,s^{-1}}} \sim 2.1 \cdot 10^{-6} \approx 0.4''.$$

This angle, though small, is possible to detect: The required lens diameter is roughly

$$d \sim \lambda/\theta \sim \frac{0.5 \cdot 10^{-6}\,\mathrm{m}}{2.1 \cdot 10^{-6}} \sim 20\,\mathrm{cm}.$$

The eclipse expedition of 1919, led by Arthur Eddington of Cambridge University, tried to measure exactly this effect.

For many years – between 1909 and 1916 – Einstein believed that a correct theory of gravity would predict the Newtonian value, which turns out to be 0.87 arcseconds for light that grazes the surface of the sun. The German mathematician Soldner derived the same result in 1803. Fortunately for Einstein's reputation, the eclipse expeditions that went to test his (and Soldner's) prediction got rained or clouded out. By the time an expedition got lucky with the weather (Eddington's in 1919), Einstein had invented

This sentence saying how to measure the bending of light seems out of place sandwiched between two sentences on how and why we can't measure bending light on earth. Perhaps it should be relocated to the end of the next paragraph, where you actually talk about this?

I like this position because it helps the reader understand the use for much of the math that is being used. Moving this sentence to the next paragraph would make the last sentence of this paragraph difficult to understand.

Interesting thing I found out just now from wikipedia: to be completed in 2018, the Giant Magellan Telescope will have 7 8.4m mirrors that, together, have the resolving power of a single 24.5m mirror. Also, for reference, the hubble has a 2.5m mirror.

can you show how the two equations for the angle would relate to each other?

is there a way to magnify the result of a small lens to somehow emulate the result of a d 500m lens?

Well, one could use a much shorter wavelength of light, since the minimal angle is lambda/d

This is cool to think about. If one explanation does fit, find another!

I like how from here on out you include some history...it's interesting and very relevant to the section!

why did they choose the largest mass? ...it seems like the densest planet would be a better choice (largest mass with smallest radius)...but I'm not sure.

Wow this surprised me. I didn't think that the light from the Sun would be bent. Also, would using a comparison to slope help explain why light doesn't bend?

It might be nice to note that this is in arcseconds.

That's not that small...aren't most older telescopes this size anyways?

Possibly, but this implies a perfect lens

So even here the value of the angle is really small.

I've never heard of him.

$$\frac{Gm}{rc^2} \sim \frac{6.7 \cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 6.0 \cdot 10^{24}\,\mathrm{kg}}{6.4 \cdot 10^{6}\,\mathrm{m} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}}} \sim 10^{-9}.$$

This miniscule value is the bending angle (in radians). If physicists want to show that light bends, they had better look beyond the earth! That statement is based on another piece of dimensional analysis and physical reasoning, whose result is quoted without proof: A telescope with mirror of diameter d can resolve angles roughly as small as λ/d, where λ is the wavelength of light; this result is based on the same physics as the diffraction pattern on a CDROM (Section 1.1). One way to measure the bending of light is to measure the change in position of the stars. A lens that could resolve an angle of 10^{-9} has a diameter of at least

$$d \sim \lambda/\theta \sim \frac{0.5 \cdot 10^{-6}\,\mathrm{m}}{10^{-9}} \sim 500\,\mathrm{m}.$$

Large lenses warp and crack; one of the largest existing lenses has $d \sim 6\,\mathrm{m}$. No practical mirror can have $d \sim 500\,\mathrm{m}$, and there is no chance of detecting a deflection angle of 10^{-9}.

Physicists therefore searched for another source of light bending. In the solar system, the largest mass is the sun. At the surface of the sun, the field strength is

$$\frac{Gm}{rc^2} \sim \frac{6.7 \cdot 10^{-11}\,\mathrm{m^3\,s^{-2}\,kg^{-1}} \times 2.0 \cdot 10^{30}\,\mathrm{kg}}{7.0 \cdot 10^{8}\,\mathrm{m} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}} \times 3.0 \cdot 10^{8}\,\mathrm{m\,s^{-1}}} \sim 2.1 \cdot 10^{-6} \approx 0.4''.$$

This angle, though small, is possible to detect: The required lens diameter is roughly

$$d \sim \lambda/\theta \sim \frac{0.5 \cdot 10^{-6}\,\mathrm{m}}{2.1 \cdot 10^{-6}} \sim 20\,\mathrm{cm}.$$

The eclipse expedition of 1919, led by Arthur Eddington of Cambridge University, tried to measure exactly this effect.

For many years – between 1909 and 1916 – Einstein believed that a correct theory of gravity would predict the Newtonian value, which turns out to be 0.87 arcseconds for light that grazes the surface of the sun. The German mathematician Soldner derived the same result in 1803. Fortunately for Einstein's reputation, the eclipse expeditions that went to test his (and Soldner's) prediction got rained or clouded out. By the time an expedition got lucky with the weather (Eddington's in 1919), Einstein had invented

I'm a bit confused on how an eclipse might prove the result?

> I think it was so they could look at the stars behind the sun, which would not normally be visible due to the suns brightness, but I'm not 100% sure.

It's really interesting looking at the history going on a this time as well. According to 8.225, this was one of several efforts to measure this. I think the first one failed because the person sent to measure the eclipse got caught in the war and couldn't actually make the measurements until much later.

If you can insert the images of the starfield and the starfield during the eclipse, that would be really helpful here.

how do you get arcseconds?

> An arcsecond is just 1/60 of a degree.

It's interesting to think about how much luck/chance has in scientific discovery.

wow. weather is so important. this is why we should get pretty days all the time! :)

a new theory of gravity – general relativity – and it predicted a deflection of 1.75 arcseconds.

The goal of Eddington's expedition was to decide between the Newtonian and general relativity values. The measurements are difficult, and the results were not accurate enough to decide which theory was right. But 1919 was the first year after the World War in which Germany and Britain had fought each other almost to oblivion. A theory invented by a German, confirmed by an Englishman (from Newton's university, no less) – such a picture was comforting after the trauma of war. The world press and scientific community saw what they wanted to: Einstein vindicated!

A proper confirmation of Einstein's prediction came only with the advent of radio astronomy, in whichsmall deflections could be measured accurately. Here is then a puzzle: If the accuracy (resolving power) of a telescope is λ/d, where λ is the wavelength and d is the telescope's diamater, how could radio telescopes be more accurate than optical ones, since radio waves have a much longer wavelength than light?

It's funny how life works out.

I like the history lesson of the discovery. It makes it more interesting.

i like the history lesson actually

agreed - it's interesting

I'm not sure I understand why that particular piece of information is relevant. It's interesting, of course, but does it somehow make the Englishman more English?

As scientists taught by scientists, I don't think our inclination is to believe that non-scientific forces can affect scientific outcome. I agree that I like the history anecdote.

This is an interesting comment.

that is really funny!

minor grammar? either take out the "to" or add a word, such as "see" after

(typo)

I don't understand this completely. How can you measure light diffraction off a planet or the sun with radio waves?

a new theory of gravity – general relativity – and it predicted a deflection of 1.75 arcseconds.

The goal of Eddington's expedition was to decide between the Newtonian and general relativity values. The measurements are difficult, and the results were not accurate enough to decide which theory was right. But 1919 was the first year after the World War in which Germany and Britain had fought each other almost to oblivion. A theory invented by a German, confirmed by an Englishman (from Newton's university, no less) – such a picture was comforting after the trauma of war. The world press and scientific community saw what they wanted to: Einstein vindicated!

A proper confirmation of Einstein's prediction came only with the advent of radio astronomy, in whichsmall deflections could be measured accurately. Here is then a puzzle: If the accuracy (resolving power) of a telescope is λ/d, where λ is the wavelength and d is the telescope's diamater, how could radio telescopes be more accurate than optical ones, since radio waves have a much longer wavelength than light?

Does this have to do with the fact that arrays of radio telescopes can have effective diameters that are incredibly large?

It could be due to the enormous size of radio telescopes (unless I'm confused on what rad teles are) compared to optical ones. Alternately, radio waves may be less susceptible to distortion in ways that would make them more useful / accurate over long distances?

Yep; it's called interferometry, so you can have radio telescopes with an effective diameter the length of the Earth, so it winds up having a higher resolution.

I like how throughout this class we've been realizing why all these little things are the way they are. I never really would have questioned how they determined that particular dimension, but now we know.

I just visited the largest radiotelescope in the world, and seeing this in this week's reading was fantastic!

It also helps that lamda_radio/lamda_visible_light is between 10ˆ7 and 10ˆ9

I also visited the largest radio telescope this past week. The dish itself had a diameter of 300m because you don't have to worry about lenses. So the ability to have a much larger base gives a better resolution.

sweet, I was just reading about this the other day!

VLA?

I like it when you wrap up each section with a concluding paragraph or a take-home message about how to solve these problems. Especially since these reading was a bit longer than usual, would it be possible to add a conclusion to recap what we learned in this section? I'm feeling a bit scattered right now.

I agree. Even though we're currently in a history-laden section, some sort of recap of the principals/methods discussed would be helpful - at least to help solidify the important points.

Yeah, I was expecting something different with the last few paragraphs.

I'm alright with it as long as in the next section there is some explanation of the puzzle. I think the history is alright after all the calculation explanation, and this ending just leads me to believe the section isn't over. I think this may have been a bad place to end a reading, since I'm left to want to know more (and am slightly confused since I don't know the answer to the end).

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of **independent** dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

Comments on page 1

Read Section 5.5 (Buckingham Pi) and 5.6 (drag) for the next memo. It's due Wednesday at 9am (since this is posted later than usual). I've also included the end-of-chapter problems, which is why this assignment seems longer than usual, but there's no need to comment on those pages.

I am not too happy with the transition from 5.4 to 5.5. Other sections seemed to have a progression, but this transition seems to me very distant and leaves the previous section hanging. It feels better to read them at very separate times than one after the other.

I thought that this entire section was very clear and a good way to explain the theorem. It may be useful to go back and show the actual theorem at the end of this section though.

I agree... i really like how he relates it to an example we recently went over as well. It greatly helps the understanding.

...step in any dimensional...

typo

insert: made

This?

Do you mean "That task is made simpler..."?

What are we referring to when we say groups

the B. Pi theorem may even be useful a little earlier – we learn this in 2.006, and so I was already thinking about it when reading the previous sections.

Is this somehow related to Buckingham Palace?

haha, unfortunately it doesn't. The theorem is named after a late 19th - early 20th century physicist.

http://en.wikipedia.org/wiki/Buckingham_%CF%80_theorem

It is important for this topic, specifically dimensional analysis. ;)

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of* ***independent*** *dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

Not familiar with this...maybe some background?

It's not clear here, but he's about to describe how to do it. It might be helpful to mention that it's the method used to form dimensionless (Pi) groups.

I agree, even a sentence here, or maybe for the intro, about what the buckingham pi theorem is, would be nice!

Oh that helped - i didn't have any clue how pi was related to the discussion

Is there a reason for it to be named Buckingham?

This is good because it lets us know why we're doing all of the example that we do later.

i don't exactly like the way this is introduced. "possible beginning" sounds a bit awk

why is this helpful to us?

It helps you figure out when you are done coming up with groups

how about rewording it as: "it is helpful to begin your theorem statement as 'the number of dimensionless groups is...' "?

does this theorem only help you decide after you pick out the relevant parameters. because if that is the case it doesn't seem like it would be that useful

It would be helpful for me to see a table like the ones before with the dimensions and meanings of each variable (theta, G, m, r, c) just so that I could see concretely that these groups are dimensionless and then be able to more easily find additional examples myself.

I think this list of dimensionless groups all stems from theta and Gm/rc2. Everything in this list is multiples of those, so a table wouldn't really do much clarification.

I agree, the paragraph may seem a bit mathematical, but having prior knowledge of what is being talked about makes it understandable.

Why is this include as well as theta? doesn't this lead room for theta^3?

I think the point here is that there are infinite possibilities with just theta alone, since it is dimensionless and thus any theta (squared, cubed...and on) would also be dimensionless

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is. . . .* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of* ***independent*** *dimensionless groups is. . . .* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

How did we jump from light bending to hydrogen? Maybe you shouldn't mention the light bending problem first if it can't be solved

We're just going over the results we've derived before, so I think this is fine in context.

We are just briefly revisiting problems we've seen in previous sections

I like that he went back to the previous section, though i can see why the jump would be confusing. maybe open with this and then bring on the light bending?

I am not entirely confident on what is being explained here. Is it simply a motivation for why 'independent dimensionless groups' is a better statement?

It's showing that if we don't require independence, then the number of dimensionless groups is either 0 or infinity, which isn't a very useful conclusion.

Is he trying to tease us of what the theorem concludes? He explains it's possible to have infinite or zero groups.

It makes sense first read-through. I think he's just elaborating on an example we already looked at.

Yeah, the main idea is at the last sentence?

Is this saying that for ANY list, the number of possibilities for dimensionless groups is either 0 or infinity?

Yes.

It says it seems that way, it doesn't feel like thats it though.

It is true... you can always take a different multiple or exponent and have a different dimensionless group. The key here which is explained in the next paragraph is the idea of independent dimensionless groups

You are right, although we don't care about multiples...just exponentials.

I think the statement is accurate...it does seem that a dimensionless group could be formed from almost any combination of variables once you get started...but then maybe you start to think about the number of combinations possible considering the different variables involved.

This is logical though, because if there is a dimensionless group at all, you can manipulate any number of dimensions by multiplication/cancelleation of choice to maintain that.

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is. . . .* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of **independent** dimensionless groups is. . . .* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

So if the amount of dimensionless groups is zero how can we form a relation and solve the problem?

By what you showed earlier it seems like if one group is possible then infinite dimensionless groups are possible?

Yup. Like with theta, and any combo that ends up dimensionless...since now there are no dimensions, you can divide it by any unity, or just square it or cube it or modify it in any way you want to so long as you don't add an extra diimension

I like the way that you did this. It draws attention to the fact that it is important that the groups must be independent without just saying "this is important." I think doing it this way makes it easier to understand (things are broken into pieces) and helps you remember it better

Seeing how "independent groups" is used frequently throughout both of these sections it might be helpful to explicitly explain what they are.

I think it just means that the group isn't a multiple of a group you already came up with or the addition of two groups...

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is. . . .* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of* ***independent*** *dimensionless groups is. . . .* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

in that they have no overlapping variables?

I think so

I think they can have overlapping variables...I think this means that one dimensionless group is not a function of the other. For example, theta and Gm/rcˆ2 would be independent, but theta and thetaˆ2 would not be.

Nope - actually you guys are right, no overlapping variables. I realize that that's exactly the same thing I said (not functions of each other) so I basically contradicted myself...

Here independent means that you can't construct one group from some combination of the other groups.

In any case, it should probably be explicit in the reading.

you can't construct one from a combination of the others is the correct definition of "independent." It says nothing about each variable on its own, because on their own they are likely not dimensionless.

i don't understand what 10:37 means. It seems just that one can only use a certain variable once in a group.

Right, but it also means that if you are going to set them equal, you cant have the same variable on both sides (they would cancel) so if X is used in XR/L, then it can't also be used in XYG because setting them equal would make the Xs cancel and no longer be dimensionless (I chose the letters at random, so they don't mean anything)

I thought that this little lead in didn't read as smoothly as you had intended. Perhaps changing the order of the sentences and ending with "The number of ... is" would be better? It just seemed kind of clunky.

i think that because this is more familiar to you, you can just <know> which are convenient choices. this is much more difficult for someone like me, who is just starting out.

So what would be a different set of independent groups we could create? I'm only seeing the obvious one that we chose.

I agree... I don't see how there is room for improvement or change...

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of* ***independent*** *dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

I'm a little confused by this statement. If there are only 2 independent groups, how can the two choices be "convenient?" If they are the only ones, they are the only ones (unless I'm misunderstanding something here).

I think he might be saying that we could have chosen the two groups differently (a different set of two dimensionless groups), but that our choice of having one group as theta was convenient for what we were trying to accomplish in the problem.

I agree. Also, doesn't theta have immediate dimensions?

I thought in class he has said that theta is dimensionsless

I believe we're using theta in radians not degrees, in which case it is dimensionless.

Yeah, theta is dimensionless. it has units, but it's still dimensionless. there is a difference between units and dimensions. he went over this in class.

Until he went over this in class I was a little confused - they might've mentioned it in other classes but it definitely didnt stick.

In general, angles are dimensionless. Think of all those times you've approximated sin(theta) as theta. sin definitely doesn't have dimensions, so entiehr does theta

So I like the idea of dimensional analysis, but it seems like a pain to go through the units yourself trying to cancel them all out. How does this theorem make this easier?

I really like how you refer back to some problems we saw in the previous section. Also its helpful that you put them in a list instead of paragraph form because it makes its easy to read through them and remember what the results from before were.

What does this size correspond to physically?

This is the Bohr radius between the center and the single orbiting electron.

So is it just h that makes the difference?

not exactly happy with the transition between 2 and 3. i think it's an excellent example but i think there could be some sort of transition so we know where you are going would be helpful. i found myself rereading 2 a few times, had i known that we'd find out in 3, i'd understand why you did 2 better without getting stuck in between

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of **independent** dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

Maybe list the independent group here, since the groups were listed above for light bending.

what was h again?

I believe it represents the atomic number because we are using a0 as the size. In this case it would just be 1.

would be nice to clarify what that group is just for the record.

is it always 3 or is it the number of dimensions. there are others besides length, mass, and time. temperature is one for example

# of dimensions. If there was something with length, time, mass, and temperature, you would subtract 4 for example

I think you probably made this comment before reading the next page, but it shows that your thinking is right on track!

wow this is so cool, but is it always true?

I have a vague memory of being taught that once is an incident, twice is a coincidence and thrice is a pattern. Is this true?

I think you can have a pattern if you give the rule ahead of time. If we were just empirically observing things, we couldn't say anything about 1 or 2 observations, but we might be able to figure out a pattern with 3 or more observations (which speaks to your point). However, since we have the rule ahead of time, we can call it a pattern right off the bat.

I think that's more an adage than a hard rule. 3 examples are often used to show something, but as a rhetorical device, not a scientific one.

where did this number come from?

This 3 just came from looking at the three examples he just listed...they all have 3 more variables than independent groups.

It almost seems too easy, though - usually a pattern is more complicated...

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of* ***independent*** *dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

what are quantities? variables?

What else could they be? We're experimenting with dimensionless groups.

I prefer the # to no. notation here but thats just personal preference.

agreed when I first saw this I thought you were saying no to negate previous statement.

I thought we had endless possibilities with the bending of light quantities? Following this, we get 2 possibilities (5quantities -3). im confused

i concur to this being distressing. are there no counter examples at all?

Even more distressing, I think, because you start off with negative 3 pi groups.

I'm confused by this post. When are we multiplying something by pi?

I believe this post is referring to it as a Pi group, not 3*pi groups. Meaning, before we do any analysis, we start with negative 3 independent groups (or pi groups).

Yeah, that is in fact what I meant. I guess it's not stated in the text, but these independent groups are called pi groups, hence the Buckingham Pi Theorem. The use of pi as the representative symbol for the dimensionless groups comes from Buckingham's use of the symbol when he was first writing about them in 1914.

If I recall correctly, it had something to do with taking products (sort of like how you use sigma for summation, you use pi for adding up multiplications). Nothing to do with 3.14

Hmm. I think a line about how Pi is used to describe the groups would be well placed in this introduction. I think it makes the idea of the Buckingham Pi theorem a little easier to follow.

First, I thought you meant the 3 listed above.

well, he does. as a rule, but not the three in the list

...so basically dimensional analysis is pattern matching. Please correct me if I'm wrong.

It is pretty clear when I read the rest of the sentence, but I think it is a little distracting.

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of* ***independent*** *dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions}.$$

Test this statement with additional examples:

so this is the number of basic dimensions that make up the quantities we are examining?

I believe so. I can't quickly think of anything that doesn't involve some combination of these 3 dimensions.

there are others like candela (like luminous intensity), temperature, etc. I think there are a total of 7 basic dimensions in SI however, these are the most common ones and the only ones relevant to our current analysis.

Isn't it common to have something that only involves two of these quantities though?

I also don't understand how the other dimensions such as charge are not factored into this theorem. Are they somehow not as fundamental as Length, Mass, and Time? Or are they included in a larger more comprehensive theorem?

Sanjay, you should specify whether this state meant is only for this example or for all examples.

I've seen the Buckingham Pi theorem before (2.006), so perhaps I'm not a fair judge, but this all seems clear and well written to me.

So is the number always 3 or does it change with the number of dimensions? What if we added charge? You might go into this below, since I haven't gotten there yet, but maybe it should be explicit up here if it is always or isn't always =3. (kind of like a few...)

It changes with the number of dimensions. I think it's fairly explicit here. In the example, we are only dealing with length, mass, and time. Perhaps if we added angle or some other dimension we would see a change to the pattern.

i've never seen this before it's really cool

Although it's a review for me from 2.006, it's presented a little differently here which helps me understand the material even better.

5.5 Buckingham Pi theorem

The second step is in a dimensional analysis is to make dimensionless groups. That task is simpler by knowing in advance how many groups to look for. The Buckingham Pi theorem provides that number. I derive it with a series of examples.

Here is a possible beginning of the theorem statement: *The number of dimensionless groups is....* Try it on the light-bending example. How many groups can the variables θ, G, m, r, and c produce? The possibilities include θ, θ^2, Gm/rc^2, $\theta Gm/rc^2$, and so on. The possibilities are infinite! Now apply the theorem statement to estimating the size of hydrogen, before including quantum mechanics in the list of variables. That list is a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e. That list produces no dimensionless groups. So it seems that the number of groups would be zero – if no groups are possible – or infinity, if even one group is possible.

Here is an improved theorem statement taking account of the redundancy: *The number of* ***independent*** *dimensionless groups is....* To complete the statement, try a few examples:

1. Bending of light. The five quantities θ, G, m, r, and c produce two independent groups. A convenient choice for the two groups is θ and Gm/rc^2, but any other independent set is equally valid, even if not as intuitive.
2. Size of hydrogen without quantum mechanics. The three quantities a_0 (the size), $e^2/4\pi\epsilon_0$, and m_e produce zero groups.
3. Size of hydrogen with quantum mechanics. The four quantities a_0 (the size), $e^2/4\pi\epsilon_0$, m_e, and $\hbar$ produce one independent group.

These examples fit a simple pattern:

$$\text{no. of independent groups} = \text{no. of quantities} - 3.$$

The 3 is a bit distressing because it is a magic number with no explanation. It is also the number of basic dimensions: length, mass, and time. So perhaps the statement is

$$\text{no. of independent groups} = \text{no. of quantities} - \text{no. of dimensions.}$$

Test this statement with additional examples:

How did we determine this was the number of dimensions?

I think this just refers to the basic dimensions in the sentence above (length, mass, time).

It seems this number would be the number of dimensions included in the quantities used...which like you said is length mass and time in the previous examples

I think that, at this point, it is simply a hypothesis based on the observation that there are (almost) always 3 basic dimensions that apply to most problems: length, mass, and time.

Does this have anything to do with 3D objects and the equation relating faces, edges, and vertices?

that doesn't feel right, but i'm not sure

Are we still using the minimum amount of quantities required in order to form a dimensionless group? (like in class we added c in order to get a dimensionless group). Without defining this it seems like there would be a lot of variability

is the number of dimensions the only thing that this can represent? It seems like that was just a guess.

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

$$\text{\# of independent groups} = \text{\# of quantities} - \text{\# of } \textit{independent} \text{ dimensions}.$$

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{drag} \sim \rho v^2 A, \tag{5.1}$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

The auto switching to the current page is a very nice upgrade

oh, now I see what you mean by quantities.

I knew what "quantities" meant from the start, but I can understand why someone could be confused. Perhaps it would be more clear if the variables were more frequently introduced as quantities throughout the chapter.

Since T and x_0 are explained in one word each, it might be more coherent to do the same for k- spring constant and m-mass

i think he assumed that we would know what those meant while T and X_0 can be ambiguous for time and distance. But this probably isn't a very good assumption for readers

Time and distance could refer to different measurements in this setup, but mass and spring constant are extremely straightforward. It's a perfectly valid assumption.

so when designing dimensional analysis problems should we decide the quantities based on our intuition?

These are the variables used in a spring mass system relating to the dimensions described above. I don't think it's intuition as much as it is prior knowledge.

I have been having trouble figuring out what parameters I can use to determine the dimensionless group. How do u choose these parameters? There seems to be some a priori knowledge about the dimensionless group that allows us to find the parameters, but trhat knowledge doesnt seem evident.

maybe you also wanna say that there are three dimensions (to be absolutely clear)

well nevermind my earlier comment about being absolutely clear by including dimensions. This does a good job of clearing things up.

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

$$\#\text{ of independent groups} = \#\text{ of quantities} - \#\text{ of } independent \text{ dimensions}.$$

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{drag} \sim \rho v^2 A, \quad (5.1)$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

So would a different way to state the theorem be: # of variables - # of dimensions that make up those variables? In this case, 3-2=1.

I think that's exactly how we defined it above.

that helped clear things up for me a bit...thanks!

I was a bit skeptical at first, since 3 coincidentally happened to be the number of dimensions, but it's really cool that this actually works.

agreed

I like the way you phrased this

Always causing trouble. But I liked this example and clear derivation!

have you considered rewriting this as like 3-3=0.

"However" and maybe add "so there is at least one." to the end of the sentence.

The word is "dimensions" right?

Why are these in a different color? or is it just my laptop?

they're the same color on my laptop

Great table!

well explained. It would seem to me now that we should always look to use as many independent, as opposed to duplicate, pairs in any approximation. Should we only worry about independents here, or look for them in other situations as well?

This is a great explanation and really helpful in understanding this theorem.

Agreed. this section has been pretty good about explaining things as questions arise

Is it correct to call LT-2 a dimension? I am beginning to become confused about the difference between Dimensions, units, and just regular quantities. What exactly is LT-2?

I'm confused why LT^-2 can be used as a single dimension. Is it because both L and T only appear as LT^-2 and no where else?

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

$$\text{\# of independent groups} = \text{\# of quantities} - \text{\# of } \textit{independent} \text{ dimensions}.$$

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{drag} \sim \rho v^2 A, \quad (5.1)$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

I like this part! very clear. i think people could easily mistaken that as 3 independent groups instead of 2

Yeah I would have definitely used 3, but this makes more sense now that I see it.

This is very clearly explained.

super clear...gets the point right across

I also found this sentence helpful in clarifying the independence point raised before.

I really liked this part as well. I feel that the example was clear and familiar enough for me to understand the theorem.

I'm confused. Why can we say that LT^-2 can be treated as a single dimension? I would first think that L and T themselves are dimensions. Again, the big question I've been getting reading this is the definition of independent.

This is a really excellent derivation of an interesting result!

I still feel like this was pulled out of a hat, yes we've gone from 3 to the number of independent dimensions, but is there some way to motivate why the Buckingham Pi theorem is true, without just examining a number of examples?

I agree. I don't really feel like this is at all a proof, even if we are being incredibly liberal in our definition of proof. I mean, you could have speculated that the 3 means anything, say "a few". Of course we would have found out a few examples later that this doesn't work, but how can you better motivate that the number 3 really means the number of dimensions?

Are there some other analogues to this sort of equation? I feel like it's sort of like in geometry where the number of vertices and edges are similarly related by a fixed difference. Or solving sets of linear equations, where you need n independent equations to solve for n variables.

yeah, I feel like we've used very specific cases to narrow things down to the theorem, and now we're using it to generalize everything, which seems a bit wrong...

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

# of independent groups = # of quantities−# of *independent* dimensions.

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{\text{drag}} \sim \rho v^2 A, \tag{5.1}$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

Ok, this is a nice and interesting formula. But I'd really like an explanation of why it's useful. It feels like it's thrown in here, and it doesn't feel connected with the earlier readings on dimensional analysis.

I think it's pretty obvious why this might be useful and I think it will be shown in the next section. I'm curious why this formula wasn't brought up earlier in the chapter.

I felt like this was what we have been trying to derive for the past few paragraphs. Also, the result of this gives us the number of dimensionless groups, the centerpiece of dimensional analysis.

This is great that this is emphasized here, read in the next section. You'll quickly realize how useful this is.

This is an interesting distinction. I would have thought to make three groups, but obviously the fact that two are related makes that incorrect. How would one determine the connections without knowing the formulas, though?

does the "independence" here has anything to do with notions of independence such as in probability or linear algebra?

I think it just means it appears alone as a dimension of a quantity. If a dimension pair appears along but the two dimensions creating the pair do not then it is an independent dimension pair.

I think the idependent needs to be defined right off. I keep flip-flopping my definition of independent from no overlapping variables to non creation out of each dimension. But actually I think I am just lost.

It's interesting that you mentioned that because I was thinking the same when I read through this (regarding references to probability and linear algebra and determining independence). At first, when I read that there are either 0 solutions or infinite solutions, I made the comparison that in linear algebra, if there is one solution to Ax = 0 or Ax = b, then there are infinite solutions since we can have linear combinations of that solution. With probability, this reference came to mind when reading this sentence exactly, looking at independent dimensions, particularly pairwise independence (and also dependence correspondingly). Where mass and force (and acceleration and force, respectively) are independent, mass, acceleration and force are pairwise dependent. At least that was my intuition.

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

$$\text{\# of independent groups} = \text{\# of quantities} - \text{\# of } \textit{independent} \text{ dimensions}.$$

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{drag} \sim \rho v^2 A, \tag{5.1}$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

where is the [3] referencing to?

There are references missing thoughout the reading–I think Sanjoy said they'll be added in a future edit.

Is this called the Pi theorem because of the 3 basic dimensions?

Yeah what does Pi have to do with it? seems like if it were actually Pi because of "3" is just a slapped on name.

I personally would have liked to see this theorem at the beginning, and then the process of explaining why. Also, is there a more formal proof? Examples just don't seem as concrete

agreed. I'd like to know what I'm trying to convince myself of through use of the examples.

No offense to this commenter, but I think getting away from what you're talking about is exactly what this class is all about. To us, the proof doesn't matter if we know it works.

I actually liked the way the text derived the theorem with all these examples, but I do think it would be helpful to have a little bit of background after we arrive at the theorem.

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

$$\text{\# of independent groups} = \text{\# of quantities} - \text{\# of } \textit{independent} \text{ dimensions}.$$

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{drag} \sim \rho v^2 A, \quad (5.1)$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

I thought this was a great way to explain the above conclusion, but I was under the impression that this was not the actual conclusion of the Buckingham Pi theorem, just a useful bit of pre-information.

I thought the actual Buckingham Pi theorem was pi1=f(pi2, ..., pi_k), or equivalently F(pi1, ..., pi_k)=0.

In fact, going back through the readings a bit, it seems like we use this fact (and continue to in the next section) without actually stating it explicitly as any sort of theorem. (I could have missed it, though.) I think fully fleshing out the Buckingham Pi theorem might make it a more readily available tool.

Where did the theorem get its name from? I also agree that a more fleshed out definition would be useful.

Edgar Buckingham, US physicist. He published papers on the subject in 1914-15.

http://en.wikipedia.org/wiki/Edgar_Buckingham

and

http://en.wikipedia.org/wiki/Buckingham_π_theorem#Original_sources

This is a great link that clears up a lot of questions. Thanks.

I like how the Buckingham Pi Theorem was mentioned and slowly built up before directly saying its formulaic meaning.

Maybe include a sentence now on what this theorem can be used for, or something like that...?

I agree, why does this theorem help us? I saw how having dimensionless groups helped us solve problems, but how does knowing the number of independent groups via the Buckingham Pi theorem help us in this course?

Is it because it helps us check if we forgot any of the dimensionless groups?

I feel like normally you'd precede or follow this with the history of the theorem.

Yes, who was this Buckingham, and why Pi?

1. Period of a spring–mass system. The quantities are T (the period), k, m, and x_0 (the amplitude). These four quantities form one independent dimensionless group, which could be kT^2/m. This result is consistent with the proposed theorem.
2. Period of a spring–mass system (without x_0). Since the amplitude x_0 does not affect the period, the quantities could have been T (the period), k, and m. These three quantities form one independent dimensionless group, which again could be kT^2/m. This result is also consistent with the proposed theorem, since T, k, and m contain only two dimensions (mass and time).

The theorem is safe until we try to derive Newton's second law. The force F depends on mass m and acceleration a. Those three quantities contain three dimensions – mass, length, and time. Three minus three is zero, so the proposed theorem predicts zero independent dimensionless groups. Whereas $F = ma$ tells me that F/ma is a dimensionless group.

This problem can be fixed by adding one word to the statement. Look at the dimensions of F, m, and a. All the dimensions – M or MLT^{-2} or LT^{-2} – can be constructed from only *two* dimensions: M and LT^{-2}. The key idea is that the original set of three dimensions are not independent, whereas the pair M and LT^{-2} are independent. So:

Var	Dim	What
F	MLT^{-2}	force
m	M	mass
a	LT^{-2}	acceleration

$$\text{\# of independent groups} = \text{\# of quantities} - \text{\# of } \textit{independent} \text{ dimensions}.$$

That statement is the Buckingham Pi theorem [3].

5.6 Drag

For the final example of dimensional analysis, we revisit the cone experiment of Section 3.3.1. That analysis, using conservation, concluded that the drag force on the cones is given by

$$F_{drag} \sim \rho v^2 A, \tag{5.1}$$

where ρ is the density of the fluid (e.g. air or water), v is the speed of the cone, and A is its cross-sectional area. What can dimensional analysis tell us about this problem?

This is all clear until maybe the last paragraph; the 'independent' dimensions confuses me, and although I still need to read on, it is definitely something that should be further explained before going to a new subsection. I'm pretty lost at the whole idea of dimensions being combined and then being independent; I feel like this then simplifies more, or is being redundant.

This was quite clear, although I do agree with people that it should've been introduced at the beginning of the explanations of deriving dimensionless groups.

It could be interesting to lack back at the end of the class to all the different methods we analyzed drag problems into one single lecture

I really like that there are a few examples that are used throughout the readings. It means we already have a familiarization with the example and can develop a better intuition.

Definitely. I also didn't quite grasp the entire background the first time through, so reinforcement also aids in my general understanding.

I like that you go back to old examples. Drag is a popular guest.

Its also really nice seeing the same problem tackled in different ways coming up with the same conclusion.

I almost want a list at the end for all the different ways we solved for drag...they are getting a bit confused in my head!

What problem? The equation?

I believe the problem is to estimate drag.

The problem is the cone experiment- I think we are supposed to discount what we already know and try to solve it using dimensional analysis and hopefully come up with the same equation that we previously found.

I think it might be more clear if it is specified as the falling cone experiment.

The strategy is to find the quantities that affect F_{drag}, find their dimensions, and then find dimensionless groups.

On what quantities does the drag depend, and what are their dimensions?

The drag force depends on four quantities: two parameters of the cone and two parameters of the fluid (air). Any dimensionless form can be built from dimensionless groups: from dimensionless products of the variables. Because any equation describing the world can be written in a dimensionless form, and any dimensionless form can be written using dimensionless groups, any equation describing the world can be written using dimensionless groups.

v	speed of the cone	LT^{-1}
r	size of the cone	L
ρ	density of air	ML^{-3}
ν	viscosity of air	L^2T^{-1}

Problem 5.3 Kepler's third law
Use dimensional analysis to derive Kepler's third law connecting the orbital period of a planet to its orbital radius (for a circular orbit).

What dimensionless groups can be constructed for the drag problem?

According to the Buckingham Pi theorem, the five quantities and three independent dimensions give rise to two independent dimensionless groups. One dimensionless group could be $F/\rho v^2 r^2$. A second group could be rv/ν. Any other dimensionless group can be constructed from these two groups (Problem 5.4), so the problem is indeed described by two independent dimensionless groups. The most general dimensionless statement is then

$$\text{one group} = f(\text{second group}), \tag{5.2}$$

where f is a still-unknown (but dimensionless) function.

Which dimensionless group belongs on the left side?

The goal is to synthesize a formula for F, and F appears only in the first group $F/\rho v^2 r^2$. With that constraint in mind, place the first group on the left side rather than wrapping it in the still-mysterious function f. With this choice, the most general statement about drag force is

Echoing the previous comment at the end of the previous page, I'm not sure why we are doing this dimensional analysis. It might be useful to give some motivation at the beginning.

He's already told us why in the previous sections when we did these examples. This is attempting to be a more math specific section.

These dimension charts are very helpful for quick glances

I like the consistency of the charts!

Force is in the other chart but not here. I immediately said in my head '4-3 = 1' but realized when I read on that it was 2. You have force in the Newton chart but not here, and I think for consistency, it should be in both.

If we didn't think of density at first, could we get to this in a roundabout way, like mass of the air & distance traveled & size of the cone?

this isn't included in the equation above? how does it fit in?

Also, this is slightly misleading because this is not the typical viscosity people denote with mu, but rather the kinematic viscosity which is the aforementioned mu/rho.

i don't really know much about fluids...what is viscosity again?

I don't understand this sentence.

Yeah I didn't understand it at first. This is a really roundabout way of stating something I think is obvious.

I do think this sentence is redundant/worded awkwardly...but it's saying that by multiplying some combination of variables, you get dimensionless products that form dimensionless groups.

Sometimes, it's confusing to put v and nu in the same place. Depending on the font, they can be almost identical...

This paragraph is kind of silly i think. It goes off on making big general statements without supporting any of them. Then it draws a conclusion from those loaded statements.

I feel like it's not that they CAN be, isn't it that they SHOULD be?

The strategy is to find the quantities that affect F_{drag}, find their dimensions, and then find dimensionless groups.

On what quantities does the drag depend, and what are their dimensions?

v	speed of the cone	LT^{-1}
r	size of the cone	L
ρ	density of air	ML^{-3}
ν	viscosity of air	L^2T^{-1}

The drag force depends on four quantities: two parameters of the cone and two parameters of the fluid (air). Any dimensionless form can be built from dimensionless groups: from dimensionless products of the variables. Because any equation describing the world can be written in a dimensionless form, and any dimensionless form can be written using dimensionless groups, any equation describing the world can be written using dimensionless groups.

Problem 5.3 Kepler's third law
Use dimensional analysis to derive Kepler's third law connecting the orbital period of a planet to its orbital radius (for a circular orbit).

What dimensionless groups can be constructed for the drag problem?

According to the Buckingham Pi theorem, the five quantities and three independent dimensions give rise to two independent dimensionless groups. One dimensionless group could be $F/\rho v^2 r^2$. A second group could be rv/ν. Any other dimensionless group can be constructed from these two groups (Problem 5.4), so the problem is indeed described by two independent dimensionless groups. The most general dimensionless statement is then

$$\text{one group} = f(\text{second group}), \tag{5.2}$$

where f is a still-unknown (but dimensionless) function.

Which dimensionless group belongs on the left side?

The goal is to synthesize a formula for F, and F appears only in the first group $F/\rho v^2 r^2$. With that constraint in mind, place the first group on the left side rather than wrapping it in the still-mysterious function f. With this choice, the most general statement about drag force is

Is it possible to give an example for this sentence? The use of the word "dimensionless" is confusing me since it it referred to 3 times. An example will clarify your point.

What is the difference between forms and groups?

I am also confused by this, But my guess would be that dimensionless form just means that it has no dimensions, while a dimensionless group is an instance of something with dimensionless form.

This is kind of cool. I never considered this before. I'm still not entirely convinced it's true, but I like this statement because it got me thinking.

I'm not convinced it's true either. It just came out of nowhere.

Actually, it has to be true. Think about it:

If I have a statment, x = y, then x and y must have the same dimensions. So let's define a constant with those dimensions called C. Thus, x/C=y/C, but x/C and y/C are dimensionless groups, so we've expressed an arbitrary relation using only dimensionless groups.

its like finding the invarient all over again

This makes a lot more sense after the comment in lecture that the world doesn't care about units/dimensions, it cares about relations.

Interesting - and I agree that the comment was quite clarifying. Dimensions, as we define them, are just relations to some known physical quantity - take, for example, the kilogram, an arbitrary block of substance. We're simply defining the relation of something else to this arbitrary baseline.

Wow. That's a pretty powerful statement. If dimensionless groups are so powerful, why is it that I've never really heard them mentioned till this class???

Yeah.. Bold statement.

The strategy is to find the quantities that affect F_{drag}, find their dimensions, and then find dimensionless groups.

▶ *On what quantities does the drag depend, and what are their dimensions?*

v	speed of the cone	LT^{-1}
r	size of the cone	L
ρ	density of air	ML^{-3}
ν	viscosity of air	L^2T^{-1}

The drag force depends on four quantities: two parameters of the cone and two parameters of the fluid (air). Any dimensionless form can be built from dimensionless groups: from dimensionless products of the variables. Because any equation describing the world can be written in a dimensionless form, and any dimensionless form can be written using dimensionless groups, any equation describing the world can be written using dimensionless groups.

Problem 5.3 Kepler's third law
Use dimensional analysis to derive Kepler's third law connecting the orbital period of a planet to its orbital radius (for a circular orbit).

▶ *What dimensionless groups can be constructed for the drag problem?*

According to the Buckingham Pi theorem, the five quantities and three independent dimensions give rise to two independent dimensionless groups. One dimensionless group could be $F/\rho v^2 r^2$. A second group could be rv/ν. Any other dimensionless group can be constructed from these two groups (Problem 5.4), so the problem is indeed described by two independent dimensionless groups. The most general dimensionless statement is then

$$\text{one group} = f(\text{second group}), \tag{5.2}$$

where f is a still-unknown (but dimensionless) function.

▶ *Which dimensionless group belongs on the left side?*

The goal is to synthesize a formula for F, and F appears only in the first group $F/\rho v^2 r^2$. With that constraint in mind, place the first group on the left side rather than wrapping it in the still-mysterious function f. With this choice, the most general statement about drag force is

Maybe you should spell out what kepler's third law is, so we know if we got it or not.

I agree! It doesn't require any explanation, just a basic equation would do.

So I was going to tell you Kepler's third, but I forgot the exponents on T and R, which I can never remember. This problem actually turned out to be really instructive. Find dimensionless groups between G, M (of star), R, and T, and figure out how T and R are related based on the powers of T and R in the dimensionless group.

I feel like this problem might be better placed if you moved it to the end of your example about drag. Here it interrupts the flow, but at the end it would give the reader an opportunity to try and apply what they had just learned in the drag example to a different problem.

I agree, the placement of Problems has been erratic before too. Putting them at the ends of sections is a good idea. One or two at the end of each section is still better than 10 at the end of the chapter, where it's easier to ignore them and harder to see how to apply what you've just read about.

I also felt that when my mind was on drag forces and the cone experiment and suddenly this example problem was thrown in, it kind of threw off my mental picture.

when this is stated so briefly and succinctly, it always throws me for a loop a little. could you provide more steps and help guide us through it?

I don't remember having problems embedded in the text in previous sections (Maybe I haven't seen one in a while). I think its pretty neat though and encourage them to be included in other chapters!

The strategy is to find the quantities that affect F_{drag}, find their dimensions, and then find dimensionless groups.

On what quantities does the drag depend, and what are their dimensions?

v	speed of the cone	LT^{-1}
r	size of the cone	L
ρ	density of air	ML^{-3}
ν	viscosity of air	L^2T^{-1}

The drag force depends on four quantities: two parameters of the cone and two parameters of the fluid (air). Any dimensionless form can be built from dimensionless groups: from dimensionless products of the variables. Because any equation describing the world can be written in a dimensionless form, and any dimensionless form can be written using dimensionless groups, any equation describing the world can be written using dimensionless groups.

Problem 5.3 Kepler's third law
Use dimensional analysis to derive Kepler's third law connecting the orbital period of a planet to its orbital radius (for a circular orbit).

What dimensionless groups can be constructed for the drag problem?

According to the Buckingham Pi theorem, the five quantities and three independent dimensions give rise to two independent dimensionless groups. One dimensionless group could be $F/\rho v^2r^2$. A second group could be rv/ν. Any other dimensionless group can be constructed from these two groups (Problem 5.4), so the problem is indeed described by two independent dimensionless groups. The most general dimensionless statement is then

$$\text{one group} = f(\text{second group}), \tag{5.2}$$

where f is a still-unknown (but dimensionless) function.

Which dimensionless group belongs on the left side?

The goal is to synthesize a formula for F, and F appears only in the first group $F/\rho v^2r^2$. With that constraint in mind, place the first group on the left side rather than wrapping it in the still-mysterious function f. With this choice, the most general statement about drag force is

It might be helpful to state that the fifth is the force itself.

I agree. I just glanced at the table and saw 4 things listed and then see the number 5 here. Maybe the force should be listed in the table too.

I don't know if it's necessary. Earlier readings mentioned that you need to consider all quantities, so the reader should remember to include force. This is presuming he/she is reading the chapter as a whole, instead of in the fragmented form they we have.

Counting on the reader to remember force seems like it's causing confusion. I think it should be included in the table.

agreed

so theoretically you can construct an infinite number of dimensionless group but the B.Pi theorem tells you only 2 of them are valid, so how do you choose the correct two from a pool of infinite groups

There are no "correct" 2. Any two independent groups will work... But it's usually easier to pick the simplest ones.

I think the B. Pi Thm says there are only 2 "independent" groups in this case, but there are infinitely many groups of 2 that would work.

Yeah, he's talking about independent groups, where they can't be divided down any further and still be dimensionless (any multiplied combination of the two is still dimensionless but not independent).

Why does F/p(vr)^2 count? This is just putting the drag force on the same side as it's equation.

This example solidified it for me.

How come these are so different than the analysis above F pv^2*A I thought that they would match more closely. do they provide similar results to each other?

The strategy is to find the quantities that affect F_{drag}, find their dimensions, and then find dimensionless groups.

▶ *On what quantities does the drag depend, and what are their dimensions?*

v	speed of the cone	LT^{-1}
r	size of the cone	L
ρ	density of air	ML^{-3}
ν	viscosity of air	L^2T^{-1}

The drag force depends on four quantities: two parameters of the cone and two parameters of the fluid (air). Any dimensionless form can be built from dimensionless groups: from dimensionless products of the variables. Because any equation describing the world can be written in a dimensionless form, and any dimensionless form can be written using dimensionless groups, any equation describing the world can be written using dimensionless groups.

Problem 5.3 Kepler's third law
Use dimensional analysis to derive Kepler's third law connecting the orbital period of a planet to its orbital radius (for a circular orbit).

▶ *What dimensionless groups can be constructed for the drag problem?*

According to the Buckingham Pi theorem, the five quantities and three independent dimensions give rise to two independent dimensionless groups. One dimensionless group could be $F/\rho v^2 r^2$. A second group could be rv/ν. Any other dimensionless group can be constructed from these two groups (Problem 5.4), so the problem is indeed described by two independent dimensionless groups. The most general dimensionless statement is then

$$\text{one group} = f(\text{second group}), \tag{5.2}$$

where f is a still-unknown (but dimensionless) function.

▶ *Which dimensionless group belongs on the left side?*

The goal is to synthesize a formula for F, and F appears only in the first group $F/\rho v^2 r^2$. With that constraint in mind, place the first group on the left side rather than wrapping it in the still-mysterious function f. With this choice, the most general statement about drag force is

can variables be in more than 1 dimensionless group? I always thought that they could not be repeated

Yes. They are independent because one group does not "contain" another group. You could not, for example, 'factor out' the first group from the second group.

But both groups may contain the same variable.

The v and nu are nearly indistinguishable.

This threw me off as well.

Like that you introduced this at the end of class Monday!

I'm confused by this. Do you mean that any new group is a function of the two already defined? If you really do mean group 1 is a function of group 2, doesn't this mean that they're not independent?

Yeah, this statement is pretty unclear to me as well. It seems contradictory to what we already read about independent groups.

But isn't it true that any dimensionless group is a function of another dimensionless group? I thought before it discussed independent dimensions, not independent dimensionless groups.

This is a fundamental point of dimensional analysis, any non-dimensional group can be described as a function of the remaining non-dimensional groups.

The tilda versus the proportionality sign i think...

how is this so different than the proportionality stuff we did with drag before? just that the units match? but how to they really differ from each other.

this statement would have been more useful earlier in the chapter...actually, this whole section would be more useful earlier in the chapter because it walks you through the method really well

I like this break down of the key points to the problem

I like that they're posed as questions, so the reader can work the following step and check it piece by piece as they go.

The strategy is to find the quantities that affect F_{drag}, find their dimensions, and then find dimensionless groups.

▶ *On what quantities does the drag depend, and what are their dimensions?*

v	speed of the cone	LT^{-1}
r	size of the cone	L
ρ	density of air	ML^{-3}
ν	viscosity of air	L^2T^{-1}

The drag force depends on four quantities: two parameters of the cone and two parameters of the fluid (air). Any dimensionless form can be built from dimensionless groups: from dimensionless products of the variables. Because any equation describing the world can be written in a dimensionless form, and any dimensionless form can be written using dimensionless groups, any equation describing the world can be written using dimensionless groups.

Problem 5.3 Kepler's third law
Use dimensional analysis to derive Kepler's third law connecting the orbital period of a planet to its orbital radius (for a circular orbit).

▶ *What dimensionless groups can be constructed for the drag problem?*

According to the Buckingham Pi theorem, the five quantities and three independent dimensions give rise to two independent dimensionless groups. One dimensionless group could be $F/\rho v^2 r^2$. A second group could be rv/ν. Any other dimensionless group can be constructed from these two groups (Problem 5.4), so the problem is indeed described by two independent dimensionless groups. The most general dimensionless statement is then

$$\text{one group} = f(\text{second group}), \tag{5.2}$$

where f is a still-unknown (but dimensionless) function.

▶ *Which dimensionless group belongs on the left side?*

The goal is to synthesize a formula for F, and F appears only in the first group $F/\rho v^2 r^2$. With that constraint in mind, place the first group on the left side rather than wrapping it in the still-mysterious function f. With this choice, the most general statement about drag force is

I was wondering that! How do you know which to define as a function of the other?

I lost what we were looking for in the long paragraph about dimensionlessness. It would be a good idea to restate the question.

I disagree, I think this does a good job of not restating what has been said before but instead moving through the problem using the information that was just described.

Well, but only because you set up the dimensionless groups like that. So I guess you have to have in mind what you're going to be looking for when you set up the groups?

i really liked this paragraph. it asked a question i was wondering and then answered it well.

$$\frac{F}{\rho v^2 r^2} = f\left(\frac{rv}{\nu}\right). \tag{5.3}$$

The physics of the (steady-state) drag force on the cone is all contained in the dimensionless function f. However, dimensional analysis cannot tell us anything about that function. To make progress requires incorporating new knowledge. It can come from an experiment such as dropping the small and large cones, from wind-tunnel tests at various speeds, and even from physical reasoning.

Before reexamining the results of the cone experiment in dimensionless form, let's name the two dimensionless groups. The first one, $F/\rho v^2 r^2$, is traditionally written in a slightly different form:

$$\frac{F}{\frac{1}{2}\rho v^2 A}, \tag{5.4}$$

where A is the cross-sectional area of the cone. The 1/2 is an arbitrary choice, but it is the usual choice: It is convenient and is reminiscent of the 1/2 in the kinetic energy formula $mv^2/2$. Written in that way, the first dimensionless group is called the drag coefficient and is abbreviated c_d. The second group, rv/ν, is called the Reynolds number. It is traditionally written as

$$\frac{vL}{\nu}, \tag{5.5}$$

where L is the diameter of the object.

The conclusion of the dimensional analysis is then

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{5.6}$$

Now let's see how the cone experiment fits into this dimensionless framework. The experimental data was that the small and large cones fell at the same speed – roughly $1\,\text{m}\,\text{s}^{-1}$. The conclusion is that the drag force is proportional to the cross-sectional area A. Because the drag coefficient is proportional to F/A, which is the same for the small and large cones, the small and large cones have the same drag coefficient.

Their Reynolds numbers, however, are not the same. For the small cone, the diameter is $2\,\text{in} \times 0.75$ (why?), which is roughly 4 cm. The Reynolds number is

Comments on page 4

Oh wow. Didn't see this coming. So for this, we haven't actually looked up the equation for Drag force right? We just analyze the components and turn them into dimensionless parts. What happens if you need three dimensionless groups though and cannot set them equal to each other?

Why must this only work in steady state?

does this function only indicate that these two dimensionless groups can both be used to calculate the drag coefficient based on some factor? I didn't see a solution for this function below

I find it interesting how we come to the two important parameters, Cd and Re, from dimensional analysis. I would have never thought to get to them this way!

How do we come up with these though? Did you just know these similar formulas, or did you look them up?

Yeah...I can't see myself coming up with these on my own...even knowing common formulas such as kinetic energy.

The drag equation may not be intuitive but something like kinetic energy is something all people should know, much less an MIT student.

Yeah, but just because you know it, doesn't mean your intuition will tell you "this is a dimensionless group".

why? just for further insight, or will this come up again later?

Was there any reason to change the function as it was?

$$\frac{F}{\rho v^2 r^2} = f\left(\frac{rv}{\nu}\right). \quad (5.3)$$

The physics of the (steady-state) drag force on the cone is all contained in the dimensionless function f. However, dimensional analysis cannot tell us anything about that function. To make progress requires incorporating new knowledge. It can come from an experiment such as dropping the small and large cones, from wind-tunnel tests at various speeds, and even from physical reasoning.

Before reexamining the results of the cone experiment in dimensionless form, let's name the two dimensionless groups. The first one, $F/\rho v^2 r^2$, is traditionally written in a slightly different form:

$$\frac{F}{\frac{1}{2}\rho v^2 A}, \quad (5.4)$$

where A is the cross-sectional area of the cone. The 1/2 is an arbitrary choice, but it is the usual choice: It is convenient and is reminiscent of the 1/2 in the kinetic energy formula $mv^2/2$. Written in that way, the first dimensionless group is called the drag coefficient and is abbreviated c_d. The second group, rv/ν, is called the Reynolds number. It is traditionally written as

$$\frac{vL}{\nu}, \quad (5.5)$$

where L is the diameter of the object.

The conclusion of the dimensional analysis is then

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (5.6)$$

Now let's see how the cone experiment fits into this dimensionless framework. The experimental data was that the small and large cones fell at the same speed – roughly $1\,\text{m}\,\text{s}^{-1}$. The conclusion is that the drag force is proportional to the cross-sectional area A. Because the drag coefficient is proportional to F/A, which is the same for the small and large cones, the small and large cones have the same drag coefficient.

Their Reynolds numbers, however, are not the same. For the small cone, the diameter is $2\,\text{in} \times 0.75$ (why?), which is roughly 4 cm. The Reynolds number is

Even at this stage in the term I don't feel comfortable taking those leaps of faith.

I was a little confused by this too, sometimes I see the connection but am not sure when I can make the bridge like this

yeah i dont see the connection either. just a random 1/2 in an equation doesn't exactly set off any alarms.

I think his point is that the random constants dont really matter that much and like mentioned, probably are only derived accurately from testing. However, since he already knows the answer, he goes ahead and "guesses correctly"

I thought constants weren't important in dimensional analysis, they are?

The point is not that you have to do it this way, but as he says "traditionally" it is done this way. You will still get a correct answer if you don't but past research and convention suggests to use this form.

yeah, this seems like a leap for me too, but i guess you could maybe see the connection. Force and kinetic energy are related, obviously, and kinetic energy is (mv^2)/2. the term we have in the denominator is (rho*v^2)/2 *A, which is kinda similar i guess..

It is a huge leap, but I think the connection here between the rather random number 1/2 and kinetic energy is cool.

I'm confused by this statement because in past readings we have dropped the 1/2 in the kinetic energy formula (and in others also I believe) because it was irrelevant to the final answer. So here, I am a little confused why we are adding it if it is just an irrelevant number?

Perhaps reword to "Written that way"

$$\frac{F}{\rho v^2 r^2} = f\left(\frac{rv}{\nu}\right). \tag{5.3}$$

The physics of the (steady-state) drag force on the cone is all contained in the dimensionless function f. However, dimensional analysis cannot tell us anything about that function. To make progress requires incorporating new knowledge. It can come from an experiment such as dropping the small and large cones, from wind-tunnel tests at various speeds, and even from physical reasoning.

Before reexamining the results of the cone experiment in dimensionless form, let's name the two dimensionless groups. The first one, $F/\rho v^2 r^2$, is traditionally written in a slightly different form:

$$\frac{F}{\frac{1}{2}\rho v^2 A}, \tag{5.4}$$

where A is the cross-sectional area of the cone. The 1/2 is an arbitrary choice, but it is the usual choice: It is convenient and is reminiscent of the 1/2 in the kinetic energy formula $mv^2/2$. Written in that way, the first dimensionless group is called the drag coefficient and is abbreviated c_d. The second group, rv/ν, is called the Reynolds number. It is traditionally written as

$$\frac{vL}{\nu}, \tag{5.5}$$

where L is the diameter of the object.

The conclusion of the dimensional analysis is then

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{5.6}$$

Now let's see how the cone experiment fits into this dimensionless framework. The experimental data was that the small and large cones fell at the same speed – roughly $1\,\mathrm{m\,s^{-1}}$. The conclusion is that the drag force is proportional to the cross-sectional area A. Because the drag coefficient is proportional to F/A, which is the same for the small and large cones, the small and large cones have the same drag coefficient.

Their Reynolds numbers, however, are not the same. For the small cone, the diameter is 2 in × 0.75 (why?), which is roughly 4 cm. The Reynolds number is

Can we perhaps get a short explanation of what the Reynolds number is? I feel like this comes out of nowhere, and it's never fully explained.

very much agree

its a dimensionless number that compares inertial forces to viscous forces in a fluid flow. this is probably more obvious to course 2s.

Agree as well, especially since it is referenced a few times in the following paragraphs

it is used to determine when a fluid becomes turbulent.

If I recall, there is a really cool video online somewhere with different reynolds number values and what happens to perturbed fluids

you should mention that we abbreviate it as Re

you should mention that we abbreviate it as Re

you should mention that we abbreviate it as Re

you should mention that we abbreviate it as Re

i guess im more used to the rho*v*D/mu form

Note that the v on the bottom is the kinematic viscosity which is equal to mu/rho, so these to terms are the same.

This is just an unfortunate property of this font, I guess, but these nu's and v's look ridiculously similar. Is there an option for another style nu, or can it be in bold or italics throughout, or some other distinguishing feature?

I agree, I actually thought this was rv/v = r until a few lines ago.

same here! maybe you can capitlize one of the variables so as to not confuse people

I was pretty confused by this too...I wondered why they didnt cancel out and had to stare really closely at the font.

Definitely not intuitive, I would have guessed they were.

$$\frac{F}{\rho v^2 r^2} = f\left(\frac{rv}{\nu}\right). \tag{5.3}$$

The physics of the (steady-state) drag force on the cone is all contained in the dimensionless function f. However, dimensional analysis cannot tell us anything about that function. To make progress requires incorporating new knowledge. It can come from an experiment such as dropping the small and large cones, from wind-tunnel tests at various speeds, and even from physical reasoning.

Before reexamining the results of the cone experiment in dimensionless form, let's name the two dimensionless groups. The first one, $F/\rho v^2 r^2$, is traditionally written in a slightly different form:

$$\frac{F}{\frac{1}{2}\rho v^2 A}, \tag{5.4}$$

where A is the cross-sectional area of the cone. The 1/2 is an arbitrary choice, but it is the usual choice: It is convenient and is reminiscent of the 1/2 in the kinetic energy formula $mv^2/2$. Written in that way, the first dimensionless group is called the drag coefficient and is abbreviated c_d. The second group, rv/ν, is called the Reynolds number. It is traditionally written as

$$\frac{vL}{\nu}, \tag{5.5}$$

where L is the diameter of the object.

The conclusion of the dimensional analysis is then

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{5.6}$$

Now let's see how the cone experiment fits into this dimensionless framework. The experimental data was that the small and large cones fell at the same speed – roughly $1\,\text{m}\,\text{s}^{-1}$. The conclusion is that the drag force is proportional to the cross-sectional area A. Because the drag coefficient is proportional to F/A, which is the same for the small and large cones, the small and large cones have the same drag coefficient.

Their Reynolds numbers, however, are not the same. For the small cone, the diameter is $2\,\text{in} \times 0.75$ (why?), which is roughly 4 cm. The Reynolds number is

Is this a result of the 1/4 of the circle you cut out before folding into the cone?

i believe so. i think he's trying to illustrate the fact that we don't have to recalculate from the first formula, since we can scale.

what is this getting at? the dimensions are fairly straight forward...

we are comparing the two cones knowing that they fall at the same speed. he is trying to show how you get to this point, i believe, so the numbers must be brought into it.

$$\text{Re} \sim \frac{1\,\text{m}\,\text{s}^{-1} \times 0.04\,\text{m}}{1.5 \cdot 10^{-5}\,\text{m}^2\,\text{s}^{-1}}, \tag{5.7}$$

where $1\,\text{m}\,\text{s}^{-1}$ is the fall speed and $1.5 \cdot 10^{-5}\,\text{m}^2\,\text{s}^{-1}$ is the kinematic viscosity of air. Numerically, $\text{Re}_{\text{small}} \sim 2000$. For the large cone, the fall speed and viscosity are the same as for the small cone, but the diameter is twice as large, so $\text{Re}_{\text{large}} \sim 4000$. The result of the cone experiment is, in dimensionless form, that the drag coefficient is independent of Reynolds number – at least, for Reynolds numbers between 2000 and 4000.

This conclusion is valid for diverse shapes. The most extensive data on drag coefficient versus Reynolds number is for a sphere. That data is plotted logarithmically below (from *Fluid-dynamic Drag: Practical Information on Aerodynamic Drag and Hydrodynamic Resistance* by Sighard F. Hoerner):

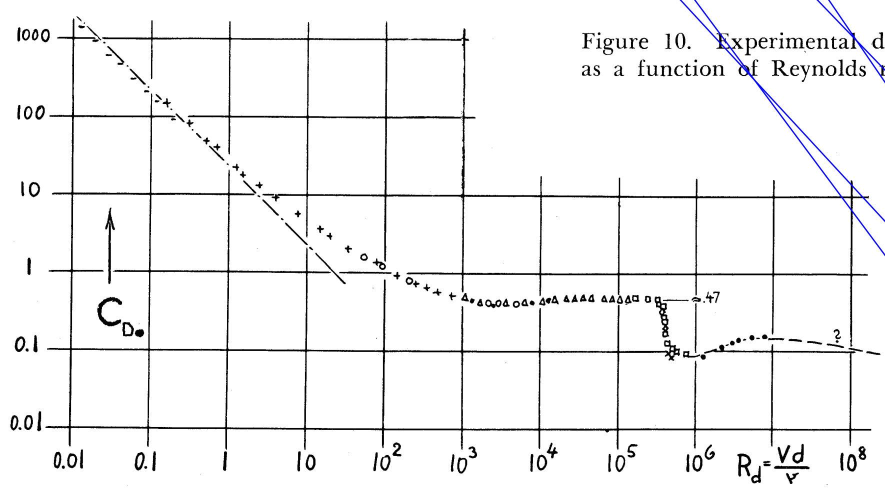

Figure 10. Experimental d as a function of Reynolds

Just like the cones, the sphere's drag coefficient is almost constant in the Reynolds number range 2000 to 4000. This full graph has interesting features. First, toward the low-Reynolds-number end, the drag coefficient increases. Second, for high Reynolds numbers, the drag coefficient stays roughly constant until $\text{Re} \sim 10^6$, where it rapidly drops by almost a factor of 5. The behavior at low Reynolds number will be explained in the chapter on easy (extreme) cases (Chapter 6). The drop in the drag coefficient, which relates to why golf balls have dimples, will be explained in the chapter on lumping (Chapter 8).

Comments on page 5

I really liked this. Now I understand the Reynold's number's used in 2.005 and a way to get them.

So does this invalidate the conclusion above that the drag coefficient can be expressed as a function of Reynolds number? Or does it simply mean that there is no easily defined function (as shown by the graph below)?

It's not that we can't express it as a function, just that, over this range, the function is a constant.

isn't the drag coefficient a function of geometry? and the cones have the same geometry so this makes sense.

I think he's trying to show that he proved this point without knowing that the drag coefficient is a function of geometry.

but didn't we just say that drag coefficient = f(reynolds number)? if it's independent, then is this not true?

Yeah this paragraph here just lost me...

yeah i'm confused

It is neat to see this insight come from an application which in the beginning of this class I would not have been able to create.

however, this indicates the switch from laminar to turbulent flow, which is important.

This is only necessarily true if we assume Cd(Re) is a monotonic function...

it'd kinda be cool to have a graph (if it exists) that has not only sphere info but also cone info (and more shapes if avail) superimposed on the same graph as a comparison

This is a good explanation, but I feel like it could have been split like the last reading into sub-sub sections (groups, reasoning, numbers) to give a bit more clarity and order.

$$Re \sim \frac{1\,\mathrm{m\,s^{-1}} \times 0.04\,\mathrm{m}}{1.5 \cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}}, \tag{5.7}$$

where $1\,\mathrm{m\,s^{-1}}$ is the fall speed and $1.5 \cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}$ is the kinematic viscosity of air. Numerically, $Re_{small} \sim 2000$. For the large cone, the fall speed and viscosity are the same as for the small cone, but the diameter is twice as large, so $Re_{large} \sim 4000$. The result of the cone experiment is, in dimensionless form, that the drag coefficient is independent of Reynolds number – at least, for Reynolds numbers between 2000 and 4000.

This conclusion is valid for diverse shapes. The most extensive data on drag coefficient versus Reynolds number is for a sphere. That data is plotted logarithmically below (from *Fluid-dynamic Drag: Practical Information on Aerodynamic Drag and Hydrodynamic Resistance* by Sighard F. Hoerner):

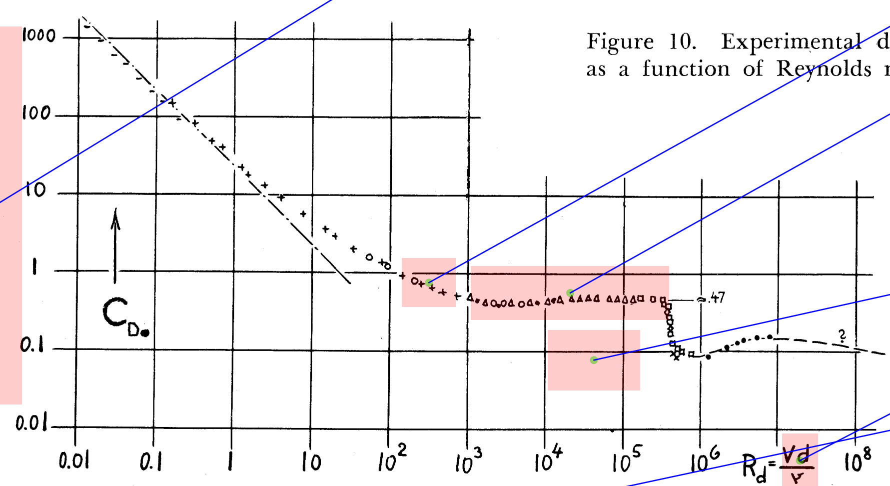

Figure 10. Experimental d
as a function of Reynolds ı

Just like the cones, the sphere's drag coefficient is almost constant in the Reynolds number range 2000 to 4000. This full graph has interesting features. First, toward the low-Reynolds-number end, the drag coefficient increases. Second, for high Reynolds numbers, the drag coefficient stays roughly constant until $Re \sim 10^6$, where it rapidly drops by almost a factor of 5. The behavior at low Reynolds number will be explained in the chapter on easy (extreme) cases (Chapter 6). The drop in the drag coefficient, which relates to why golf balls have dimples, will be explained in the chapter on lumping (Chapter 8).

I'd like to see these axes labeled

I agree, but the text does point out what is being shown and from the previous reading it should be pretty self explanatory. On another note, for a text book being published, it might be worth it to try to run these tests or find a more clear figure, this seems like it was almost hand done. Overall, the graph was extremely helpful in understand the significance of similar reynold's numbers.

Well, if you read the text, it's not that hard to figure out.

I agree! I do like that he adds real data to give insight into all this approximations.

Maybe draw a little box here to show where we are on the graph? Could be helpful to direct it for the reader so there is less wandering to understand.

Is there also an explanation for why the graph is so level at this point?

Yeah, why is it so level and then why is there a sudden drop around 10ˆ5.5?

The point of this graph, and the text, is not to explain how the Reynold's number changes, but instead to utilize the values given. Thinking too much about it's origins takes you away from the main point of the chapter.

I also wonder why it increases slightly in the middle of this constant range.

same form as listed earlier!

What does a Reynolds number within this range correspond to physically?

The point of using a dimensionless group is that you get this relationship for significantly different physical systems, you could either have a very inviscid fluid and a high velocity, or a viscous fluid and a low velocity, but provided their Reynolds number turns out the same, their drag coefficient will be the same as well.

I really like that we're doing fluids things. It brings my 2.005/2.006 studies home in an intuitive way.

"drag coefficient increases" is ambiguous. It depends if you mean as a function of moving to lower Reynolds numbers or as a standard function of Reynold number, in which case, it is decreasing.

$$\mathrm{Re} \sim \frac{1\,\mathrm{m\,s^{-1}} \times 0.04\,\mathrm{m}}{1.5 \cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}}, \qquad (5.7)$$

where $1\,\mathrm{m\,s^{-1}}$ is the fall speed and $1.5 \cdot 10^{-5}\,\mathrm{m^2\,s^{-1}}$ is the kinematic viscosity of air. Numerically, $\mathrm{Re_{small}} \sim 2000$. For the large cone, the fall speed and viscosity are the same as for the small cone, but the diameter is twice as large, so $\mathrm{Re_{large}} \sim 4000$. The result of the cone experiment is, in dimensionless form, that the drag coefficient is independent of Reynolds number – at least, for Reynolds numbers between 2000 and 4000.

This conclusion is valid for diverse shapes. The most extensive data on drag coefficient versus Reynolds number is for a sphere. That data is plotted logarithmically below (from *Fluid-dynamic Drag: Practical Information on Aerodynamic Drag and Hydrodynamic Resistance* by Sighard F. Hoerner):

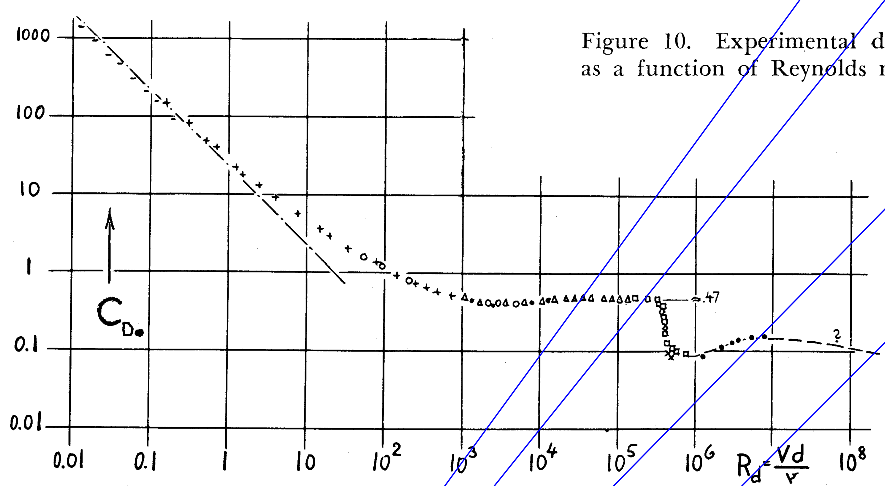

Just like the cones, the sphere's drag coefficient is almost constant in the Reynolds number range 2000 to 4000. This full graph has interesting features. First, toward the low-Reynolds-number end, the drag coefficient increases. Second, for high Reynolds numbers, the drag coefficient stays roughly constant until $\mathrm{Re} \sim 10^6$, where it rapidly drops by almost a factor of 5. The behavior at low Reynolds number will be explained in the chapter on easy (extreme) cases (Chapter 6). The drop in the drag coefficient, which relates to why golf balls have dimples, will be explained in the chapter on lumping (Chapter 8).

I thought this reading was kind of hard. I don't even know what to comment on.

It means that the laminar boundary layer becomes turbulent, thereby decreasing the amount of drag.

This is pretty interesting...thanks for dropping the term for further research.

looking forward to this discussion

Me too!

Yeah this was one of the questions off the diagnostic that I want to come out of the class knowing how to solve.

same! i was wondering when we'd get back to it

I'm glad you pointed out when we would learn these things

A comment on the whole section of dimensional analysis.... Generally, I'm pretty confused. At the end of previous sections (D&C, Abstraction, etc.), I felt like I understood the concepts and could effectively use them to solve problems. However, I'm sort of all over the place with dimensional analysis. I understood all the readings, but I just don't understand how to take these concepts and apply them to solve problems.

When we had to do pset/test problems for dimensional analysis using 2.006, the problem would go something like this: You're given a picture and description of some experiment to which you are given results, but you truly do not understand the physics. Using dimensional analysis, you form Pi groups which you know need to be functions of each other. Then you used the experimental results to determine whether the function was linear, quadratic, had some constant, etc.

Hope this helps.

Problem 5.4 Only two groups
Show that F, v, r, ρ, and ν produce only two independent dimensionless groups.

Problem 5.5 Counting dimensionless groups
How many independent dimensionless groups are there in the following sets of variables:

a. size of hydrogen including relativistic effects:

$e^2/4\pi\epsilon_0$, $\hbar$, c, a_0 (Bohr radius), m_e (electron mass).

b. period of a spring–mass system in a gravitational field:

T (period), k (spring constant), m, x_0 (amplitude), g.

c. speed at which a free-falling object hits the ground:

v, g, h (initial drop height).

d. [tricky!] weight W of an object:

W, g, m.

Problem 5.6 Integrals by dimensions
You can use dimensions to do integrals. As an example, try this integral:

$$I(\beta) = \int_{-\infty}^{\infty} e^{-\beta x^2}\,dx.$$

Which choice has correct dimensions: (a.) $\sqrt{\pi}\beta^{-1}$ (b.) $\sqrt{\pi}\beta^{-1/2}$ (c.) $\sqrt{\pi}\beta^{1/2}$ (d.) $\sqrt{\pi}\beta^{1}$

Hints:

1. The dimensions of dx are the same as the dimensions of x.
2. Pick interesting dimensions for x, such as length. (If x is dimensionless then you cannot use dimensional analysis on the integral.)

Problem 5.7 How to avoid remembering lots of constants
Many atomic problems, such as the size or binding energy of hydrogen, end up in expressions with $\hbar$, the electron mass m_e, and $e^2/4\pi\epsilon_0$, which is a nicer way to express the squared electron charge. You can avoid having to remember those constants if instead you remember these values instead:

$$\hbar c \approx 200\,\text{eV nm} = 2000\,\text{eV Å}$$
$$m_e c^2 \sim 0.5\cdot 10^6\,\text{eV}$$
$$\frac{e^2/4\pi\epsilon_0}{\hbar c} \equiv \alpha \approx \frac{1}{137} \quad \text{(fine-structure constant).}$$

Comments on page 6

Where did these problems come from?

In past sections, he always posted sample questions that relate to the reading.

Maybe from the pset we didn't have?

It seems a bit odd to have so many pages of problems. Perhaps intersperse them more throughout the chapter? The problems are as long as the material itself!

I am sure these problems will come in handy when the next pset comes out.

I thought the reading was extremely long but there are 4 full pages worth of problems - these problems are great, they're a "teaser" for what we'll see on the problem set, and get us thinking a little bit ahead of time.

it's 1, right?' mg=MG/R^2

quanties=4 number of dimensions=3

1 is the correct answer for the same reason explained above when he talked about F = ma

What exactly is this saying here? Is it some function I that is a function of B?

If [x]=L, then [B]=1/L^2, and the answer has to have dimensions of L, and the only choice with dimensions of L is b.

Could you explain why the dimensions of B are 1/L^2 if [x] = L.

B and x^2 are both in the exponent. The entire exponent term cannot have units, so everything in the power of the exponent will be a dimensionless group.

I don't see how this one is a problem; It just looks like an explanation to me. Maybe it should be in the reading?

these problems really help solidify my understanding of this technique

power radiated cannot depend on the origin. The velocity cannot matter because of relativity: You can transform to a reference frame where $v = 0$, but that change will not affect the radiation (otherwise you could distinguish a moving frame from a non-moving frame, in violation of the principle of relativity). So the acceleration a is all that's left to determine the radiated power. [This line of argument is slightly dodgy, but it works for low speeds.]

a. Using P, $q^2/4\pi\epsilon_0$, and a, how many dimensionless quantities can you form?
b. Fix the problem in the previous part by adding one quantity to the list of variables, and give a physical reason for including the quantity.
c. With the new list, use dimensionless groups to find the power radiated by an accelerating point charge. In case you are curious, the exact result contains a dimensionless factor of 2/3; dimensional analysis triumphs again!

Problem 5.10 Yield from an atomic bomb
Geoffrey Taylor, a famous Cambridge fluid mechanic, annoyed the US government by doing the following analysis. The question he answered: 'What was the yield (in kilotons of TNT) of the first atomic blast (in the New Mexico desert in 1945)?' Declassified pictures, which even had a scale bar, gave the following data on the radius of the explosion at various times:

t (ms)	R (m)
3.26	59.0
4.61	67.3
15.0	106.5
62.0	185.0

a. Use dimensional analysis to work out the relation between radius R, time t, blast energy E, and air density ρ.
b. Use the data in the table to estimate the blast energy E (in Joules).
c. Convert that energy to kilotons of TNT. One gram of TNT releases 1 kcal or roughly 4 kJ.

The actual value was 20 kilotons, a classified number when Taylor published his result ['The Formation of a Blast Wave by a Very Intense Explosion. II. The Atomic Explosion of 1945.', *Proceedings of the Royal Society of London. Series A, Mathematical and Physical* **201**(1065): 175–186 (22 March 1950)]

Problem 5.11 Atomic blast: A physical interpretation
Use energy densities and sound speeds to make a rough physical explanation of the result in the 'yield from an atomic bomb' problem.

Problem 5.12 Rolling down the plane
Four objects, made of identical steel, roll down an inclined plane without slipping. The objects are:

Just curious, if we have any extra time in class could we solve this?

It's not too bad. Give it a shot. Also, here's a link to the paper: http://tinyurl.com/y9gewzd (opens a PDF)

it's pretty interesting stuff, I just skimmed though it but it seems like a good read if you can find the time...

Oh my.

Possible answer:

E t^2 / (p R^5) = const.a

and then you can rearrange to find:

R^5 = const.b * (E/p) * t^2

where const.b is the reciprocal of const.a

this is of the form:

R = A t^(2/5)

Using the table, we fit a curve with the best guess on A.

However, here is where I run into trouble. There is still that "const." from the beginning to be dealt with.

A^5 = const.b * (E/p)

I can rearrange to calculate E:

E = const.a * p * A^5

but without any information about const.a, E could be arbitrarily large/small....

to clarify my point in a simpler/similar rephrasing:

the table only lets us initialize the RATIO of a combination of R to t, and thus we can substitute that number in to get rid of the R and t from the equation we find in part a).

But, we need some other information to initialize the arbitrary constant which our dimensionless group is equal to.

Turns out (after reading his solution), he arbitrarily set the dimensionless group to equal 1. I'm not sure why.

power radiated cannot depend on the origin. The velocity cannot matter because of relativity: You can transform to a reference frame where $v = 0$, but that change will not affect the radiation (otherwise you could distinguish a moving frame from a non-moving frame, in violation of the principle of relativity). So the acceleration a is all that's left to determine the radiated power. [This line of argument is slightly dodgy, but it works for low speeds.]

a. Using P, $q^2/4\pi\epsilon_0$, and a, how many dimensionless quantities can you form?

b. Fix the problem in the previous part by adding one quantity to the list of variables, and give a physical reason for including the quantity.

c. With the new list, use dimensionless groups to find the power radiated by an accelerating point charge. In case you are curious, the exact result contains a dimensionless factor of 2/3; dimensional analysis triumphs again!

Problem 5.10 Yield from an atomic bomb

Geoffrey Taylor, a famous Cambridge fluid mechanic, annoyed the US government by doing the following analysis. The question he answered: 'What was the yield (in kilotons of TNT) of the first atomic blast (in the New Mexico desert in 1945)?' Declassified pictures, which even had a scale bar, gave the following data on the radius of the explosion at various times:

t (ms)	R (m)
3.26	59.0
4.61	67.3
15.0	106.5
62.0	185.0

a. Use dimensional analysis to work out the relation between radius R, time t, blast energy E, and air density ρ.

b. Use the data in the table to estimate the blast energy E (in Joules).

c. Convert that energy to kilotons of TNT. One gram of TNT releases 1 kcal or roughly 4 kJ.

The actual value was 20 kilotons, a classified number when Taylor published his result ['The Formation of a Blast Wave by a Very Intense Explosion. II. The Atomic Explosion of 1945.', *Proceedings of the Royal Society of London. Series A, Mathematical and Physical* **201**(1065): 175–186 (22 March 1950)]

Problem 5.11 Atomic blast: A physical interpretation

Use energy densities and sound speeds to make a rough physical explanation of the result in the 'yield from an atomic bomb' problem.

Problem 5.12 Rolling down the plane

Four objects, made of identical steel, roll down an inclined plane without slipping. The objects are:

I always liked this problem. It causes some distress in thinking for me since it seems both intuitive and counter-intuitive at the same time. But I loved solving it in my AP physics class.

When I think of dimensions I automatically think of numbers and variables? This problem has neither....so where would i start? Would we begin assigning variables and estimated number measurement? Should I draw something or understand physical nature of the problem?

1. a large spherical shell,
2. a large disc,
3. a small solid sphere,
4. a small ring.

The large objects have three times the radius of the small objects. Rank the objects by their acceleration (highest acceleration first).

Check your results with exact calculation or with a home experiment.

Problem 5.13 Blackbody radiation

A hot object – a so-called blackbody – radiates energy, and the flux F depends on the temperature T. In this problem you derive the connection using dimensional analysis. The goal is to find F as a function of T. But you need more quantities.

a. What are the dimensions of flux?

b. What two constants of nature should be included because blackbody radiation depends on the quantum theory of radiation?

c. What constant of nature should be included because you are dealing with temperature?

d. After doing the preceding parts, you have five variables. Explain why these five variables produce one dimensionless group, and use that fact to deduce the relation between flux and temperature.

e. Look up the Stefan–Boltzmann law and compare your result to it.

What's the more specific definition of a blackbody?

If by flux you mean the energy flux density, then the dimensions are Energy/Time/Area, or M/T^3 where T is the TIME not the temperature.

If by flux you mean the energy flux density, then the dimensions are Energy/Time/Area, or M/T^3 where T is the TIME not the temperature.

If by flux you mean the energy flux density, then the dimensions are Energy/Time/Area, or M/T^3 where T is the TIME not the temperature.

If by flux you mean the energy flux density, then the dimensions are Energy/Time/Area, or M/T^3 where T is the TIME not the temperature.

reduced planck's constant h(bar) and speed of light c

Probably Kb the Boltzman constant

There are 4 dimensions, Mass, Time, Temperature and Length. 5-4=1 dimensionless group.

Flux k^4*T^4/(h^3*c^2)

Looks good! Just missing some pi's.

why so many problems?

Kinda daunting, but if you read through I think it helps. Def. dense tho.

I think that, as far as the book is concerned, yes these problems are helpful...however, to have 4 pages of problems in the reading was a little much.

That is my fault. I meant to say in the introductory NB comment, "ignore the problems,they are just there so you know what kinds of questions will be asked on the pset."

6
Easy cases

6.1 Pyramid volume 115
6.2 Atwood machine 117
6.3 Drag 120

The previous tools included methods for organizing complexity and methods for losslessly discarding complexity (for example, dimensional analysis). However, the world often throws us problems so complex – for example, almost any question in fluid mechanics – that these methods are insufficient on their own. Therefore, we now start to study methods for discarding actual complexity. With these methods, we accept a reduction in accuracy in order to reach a solution at all.

The first tool for discarding actual complexity is based on the principle that a correct solution works in all cases – including the easy ones. This principle helps us check and, more surprisingly, helps us guess solutions.

6.1 Pyramid volume

As the first example, let's explain the factor of one-third in the volume of a pyramid with a square base:

$$V = \frac{1}{3}hb^2,$$

where h is the altitude and b is the length of a side of the base.

In this reading, I thought the pyramid example was an amazing example that illustrated the concept while teaching us something really worthwhile. However, the Atwood problem I thought was a little tougher to read and was just not as awesome as the first one.

6
Easy cases

6.1 Pyramid volume 115
6.2 Atwood machine 117
6.3 Drag 120

The previous tools included methods for organizing complexity and methods for losslessly discarding complexity (for example, dimensional analysis). However, the world often throws us problems so complex – for example, almost any question in fluid mechanics – that these methods are insufficient on their own. Therefore, we now start to study methods for discarding actual complexity. With these methods, we accept a reduction in accuracy in order to reach a solution at all.

The first tool for discarding actual complexity is based on the principle that a correct solution works in all cases – including the easy ones. This principle helps us check and, more surprisingly, helps us guess solutions.

6.1 Pyramid volume

As the first example, let's explain the factor of one-third in the volume of a pyramid with a square base:

$$V = \frac{1}{3}hb^2,$$

where h is the altitude and b is the length of a side of the base.

A new tool! Read these two sections for the memo due Thursday at 10pm.

Grr, I'm sorry this is late. I pulled an all-nighter last night and then collapsed right after classes today, waking up just now.

Sounds cool! I don't know what it is.

I liked this example.

Will this section be part of a later reading?

Similarly, I think what dimensional analysis let us do was pick out the important parts.

Almost all your examples of "scary" equations come from fluids!

That and physics.

it makes me even more scared of 2.005...

well it's true and 2.005 isn't the one you should be afraid of - you don't learn much about Re or cd. You do learn it extensively in 2.006. so don't worry...yet

Does actual complexity refer to problems like those in fluid mechanics where will will have to discount more than just dimensions? I guess I just don't understand what "actual" means in this case..

I think Sanjoy might have meant something more along the lines of 'actually discarding complexity' with this sentence, however I don't really think the word "actual" is relevant to understanding the point, which is that now we're going to look at methods for discarding complexity.

agreed

What were we discarding in dimensional analysis? Virtual complexity?

Nothing was really discarded using dimensional analysis... we are moving onto a new topic (from lossless to lossy discarding of information). Although I'm not sure what information was discarded in the two examples given in the rest of the reading...

It sounds like a necessary skill set to have for complex problem solving.

6
Easy cases

6.1 Pyramid volume 115
6.2 Atwood machine 117
6.3 Drag 120

The previous tools included methods for organizing complexity and methods for losslessly discarding complexity (for example, dimensional analysis). However, the world often throws us problems so complex – for example, almost any question in fluid mechanics – that these methods are insufficient on their own. Therefore, we now start to study methods for discarding actual complexity. With these methods, we accept a reduction in accuracy in order to reach a solution at all.

The first tool for discarding actual complexity is based on the principle that a correct solution works in all cases – including the easy ones. This principle helps us check and, more surprisingly, helps us guess solutions.

6.1 Pyramid volume

As the first example, let's explain the factor of one-third in the volume of a pyramid with a square base:

$$V = \frac{1}{3}hb^2,$$

where h is the altitude and b is the length of a side of the base.

But how close would the solution be to the actual solution then? Or is this basically saying, "approximate"? I think I'm a bit confused on this.

also can you quantify the level of reduction

Could you qualify what you mean by reduction exactly? To me, estimating an answer automatically means you accept losing accuracy.

I think he's making a general statement about complex problems. He's saying that, in general, problems are often really difficult, so throwing out some of the hairier stuff will give you a solution but you will lose accuracy.

I think this is saying that what will be presented below is "lossee" estimation instead of what we've been doing so far, which is roughly "lossless estimation".

By reduction, he means he is going to "ignore" some difficult things. Say ignoring friction for something really smooth.

As we discard actual complexity, we lose accuracy by modeling the system as something that isn't identical to what is really happening

yeah, what we have been doing up to this point is approximating difficult problems but trying to model them as closely or accurately as possible. This was lossless (no critical information was lost). What he is now proposing is actively ignoring some things so that we can actually get an answer. This means that we know that we will get a wrong answer (because our model was not accurate but hopefully it will be close to the correct answer.

This principle has definitely helped me in math classes!

Yeah actually in E&M, at least the course 6 versions, we only ever use the 'easy case' of solutions. Like for solutions to the wave equation, we've only seen the easy case solution.

In other words, go with the answer that seems too obvious, right?

True, but its moreso that solving an easy case can give insight into a harder case...which really might not be solveable. It almost sounds like using an easy case (that can be solved) to roughly estimate a harder case

6

Easy cases

6.1	Pyramid volume	115
6.2	Atwood machine	117
6.3	Drag	120

The previous tools included methods for organizing complexity and methods for losslessly discarding complexity (for example, dimensional analysis). However, the world often throws us problems so complex – for example, almost any question in fluid mechanics – that these methods are insufficient on their own. Therefore, we now start to study methods for discarding actual complexity. With these methods, we accept a reduction in accuracy in order to reach a solution at all.

The first tool for discarding actual complexity is based on the principle that a correct solution works in all cases – including the easy ones. This principle helps us check and, more surprisingly, helps us guess solutions.

6.1 Pyramid volume

As the first example, let's explain the factor of one-third in the volume of a pyramid with a square base:

$$V = \frac{1}{3}hb^2,$$

where h is the altitude and b is the length of a side of the base.

I've used this idea many times without ever articulating/understanding it as a principle.

I agree. I feel like this has been discussed before in other lessons to check answers but has not been explained like this.

Same here. It's nice to have it be explained in such a way to justify what I've previously considered to be a rather hand-wave-y way to make approximations.

I'm a little confused as to what this idea is trying to convey.

I'm not sure if I understand this...do you mean the correct solution in one problem helps with others? all is it more like working backwards?

Yes I believe you work backwards by finding a solution to an easy problem and then see if it fits the general case as well.

He means that if you can't solve hard cases (like 3D) solve an easier one (like 2D) and extrapolate to more dimensions. This is similar to how we did the solitaire 'box' problem.

It's a method that you can use to solve really complex problems. First find a general solution, and then use that logic and apply it to any additional constraints, since the relationship will hold true.

I think he basically means that if you find an answer, for a problem it has to work for all cases of the problem, including the easy ones. like if you found some equation, to check it, you could plug in zero or 1 for some of the variables (simple case), and if your equation was right it would still work.

This part sounds right but I am not so sure what you really mean in context.

Yea, I had to read over this paragraph several times, but I'm pretty sure he's saying that the solution to a problem will work for simple cases as well so you can just focus on those instead of trying to check with hard or complicated examples

I think if you add a paragraph or two here with some more explanations, it might help people understand where easy cases come from. I feel that if you understand that easy cases can also be boundary conditions, or conditions where symmetry occurs, it will be easier to understand. I immediately think, ok, what are the boundaries, and then what happens at zero and infinity. Explaining this would probably be helpful.

6
Easy cases

6.1 Pyramid volume 115
6.2 Atwood machine 117
6.3 Drag 120

The previous tools included methods for organizing complexity and methods for losslessly discarding complexity (for example, dimensional analysis). However, the world often throws us problems so complex – for example, almost any question in fluid mechanics – that these methods are insufficient on their own. Therefore, we now start to study methods for discarding actual complexity. With these methods, we accept a reduction in accuracy in order to reach a solution at all.

The first tool for discarding actual complexity is based on the principle that a correct solution works in all cases – including the easy ones. This principle helps us check and, more surprisingly, helps us guess solutions.

6.1 Pyramid volume

As the first example, let's explain the factor of one-third in the volume of a pyramid with a square base:

$$V = \frac{1}{3}hb^2,$$

where h is the altitude and b is the length of a side of the base.

Good call on starting with this example. The Atwood's machine example is a big step up from this one.

I agree. This example is also very clear and easy to follow. It makes it easier to understand the technique you're trying to describe instead of just jumping into something really complex and confusing us with terminology or semantics.

A quick question about the topic of easy cases.... so the previous 2 opening paragraphs describe how we'll now begin looking at "lossy" methods of problem solving. But I'm a little confused how this pyramid volume problem illustrates lossy complexity?

I think this example as well as the Atwood example are mainly used to show how to develop and work with the "easy cases" but don't really show a lossy result. I think they're setting up for later sections with more complex problems.

I like this recognizable formula. :D

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore $1/2$ *no matter what* b *and* h *are*, so $A = bh/2$.

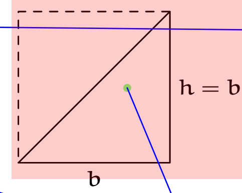

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

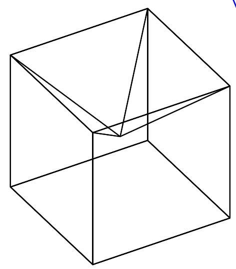

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

Hmm...I've never even thought about this problem!

I think we did this back in high school, and it was really interesting!

This is a really interesting problem, I've also never thought about it.. I always took equations like this at face value and never looked into them

I agree. This problem presented a really interesting way of looking at something I'd never really considered before.

I think we don't really need this case, it is pretty obvious and it isn't too clear how it shows the 3d case

I like the use of the 2d case. I actually never thought of a problem like this until i read through this quick example. It might be a little clearer if it was fleshed out a bit more as to what exactly we are doing, but it definitely captures the idea of finding a solution that must work simply by looking at a simple case.

yo dude speak for yourself – i like starting with baby steps.

I like this example too, I probably would be lost if you didn't build it as nicely as it is with the 2d then 3d case.

Although I appreciate why we start with easy cases, I don't immediately see the extension from 2d to 3d here.

Need we choose h = b, or could we just use a rectangle where h does not equal b, still draw a diagonal line, and still successfully show that A bh, and the constant in that is still 1/2?

I agree that any triangle would work, but if we are focusing on simple cases, setting h=b is the simplest.

Isn't this method similar to divide and conquer? I say it could be classified in that way

I think divide and conquer is taking a complex problem and breaking it down into smaller pieces. This fits more with the abstraction problems (I think)

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore $1/2$ *no matter what* b *and* h *are*, so $A = bh/2$.

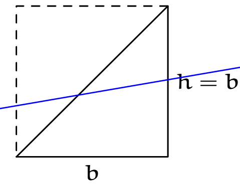

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

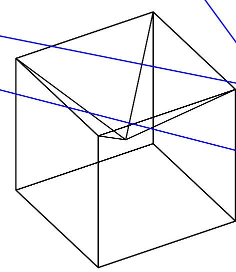

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume $4/3$. Since $hb^2 = 4$ for these pyramids, the missing constant is $1/3$. Voilà:

$$V = \frac{1}{3}hb^2.$$

Once again, this class finds me thinking about answers to questions I never would have even considered asking.

This is such an awesome example. I really like when he grabs formulas that someone made you memorize in 5th grade and you've always taken for granted and shows you where they come from.

I definitely agree with this it is awesome to reevaluate things we've taken for granted for so long and see why they really work the way they do

Also, this makes use of finding proportionality. the area of a triangle is always proportional to bh, with a constant of 1/2

Can this method be used to find the volume of different types of objects?

I like the progression from an easy example to explain a harder one.

Can't this be generalized to the case of a rectangle with two triangles that are reflections of each other?

I like informality here. Most classes I've taken focus formal proof. This one focuses on practical application.

I mean, how much formality do you need for the area of a square?

What I have liked is how visual this class has gotten... its much easier for me to understand something if I see it then if I am handed a bunch of equations.

I think part of that is also the fact that we've memorized the equations in this section like the back of our hands way back when.

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore 1/2 *no matter what* b *and* h *are*, so $A = bh/2$.

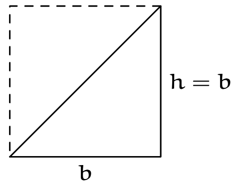

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

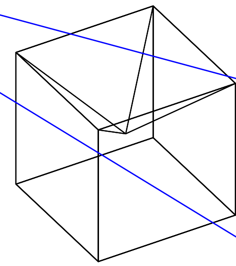

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

This may seem like a silly question, especially since everyone seems to understand it already, but it is not entirely clear to me why we can since it works for the easy case, it will work for all cases (I see how if it works for the hard cases, it works for the easy ones, but how does it work going from easy to hard?). I mean, I know what the area of a triangle is, but, I just want to make sure I know how to make the correct assumptions with harder problems.

I think it works because the simplifying assumptions - 45 degrees and b=h - make it easier to reason visually without manifesting themselves in the equations. The two assumptions are just made so it's easier to see and not so that the math is easier. For this reason it doesn't matter if we change b and h, assuming the shape (rectangle) stays the same.

why is this a major emphasis?

I'm pretty lost by this explanation...

Nevermind, unconfused. For some reason though it was the following stuff with the cube that help clear this up. I think the explanation is a bit overcomplicated which makes it a bit confusing.

This is just a minor suggestion, but you might want to include "because if it works in the easy case, it works in all cases" somewhere in here. I know you say it a few lines prior, but it might help solidify things.

I agree- this would bring it back to the main point of the topic.

I disagree, it's only been a paragraph since that statement and the italics at the end kind of reiterate that but with specifics to this problem.

Agreed, although this is common knowledge, maybe it would have been better to use a more complicated triangle to reiterate the point.

How do we know that this is not just a special case and actually is universal? We did say that, only when h=b then A=b^2.

I think what is trying to be said is that the point of these easy cases is that the principle of this estimation is to say that a solutions is valid for all.

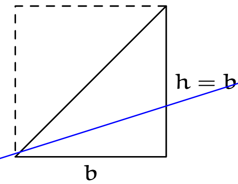

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore 1/2 *no matter what* b *and* h *are*, so $A = bh/2$.

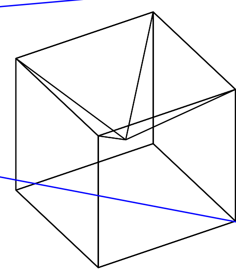

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

do you mean A/bh = 1/2 is the constant?

He wrote it correctly. He uses a twiddle and says that the constant in A bh is 1/2, so that precisely A = 1/2*(bh).

heh, it's a bit silly how much time we've spent on this whole and prop to symbols.

Could you make the mathematical connection/transition between the two cases clearer?

how can one be sure that the same principle applies in 3-dimensions? in other words, once we find the constant of 1/3 for a simple case, how can we be certain that it will apply for all hb^2 pyramids?

Extrapolation? Common sense? I dunno, I guess you would just test one or two. Or realize that the geometric properties are true regardless of the lengths of the sides

I agree with the first post here - I definitely don't intuitively grasp how the 1/2 in 2d implies a 1/3 in 3d

is this just finding the easy case? by saying the simplest case is when they combine to make a square?

I feel like this method is very case specific. Its the most intuitive process to make a shape out of the pyramid

I feel like this method is very case specific. Its the most intuitive process to make a shape out of the pyramid

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore $1/2$ *no matter what* b *and* h *are*, so $A = bh/2$.

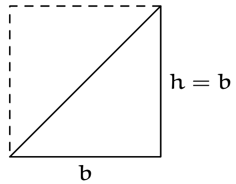

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

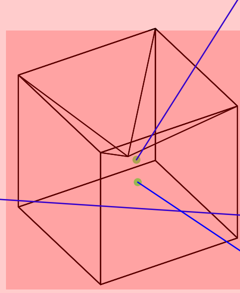

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

This is a really helpful picture for visualizing pyramids that have h =1/2 b. However, pyramids that are more extreme, like really large h but really small b, are still difficult for me to visualize forming into a cube.

genius! Coming up with the respresentation of a pyramid in a cube in this way is something that simplifies the calculation of its volume. However, how can we assume that it works for all pyramids? Is it enough to just use the triangle analogy.

this is the image i thought of originally. once you have this its easy to see how it works.

yeah this picture is super helpful, otherwise I think I would not have been able to visualize it in my head very easily

The picture is VERY helpful and VERY well done. The only suggestion that you maybe could play with to show the height better would be to maybe put a pyramid on top and on bottom to show the height of the cube is 2x the height of the pyramid.

unclear.

last time we found an easy case of h and b and made sure the area was correct, now we are doing the analogous thing in 3d.

so It's 1/6 the volume of a cube?

yes, for a cube of side length b and pyramids where b=2h, exactly 6 pyramids have the same volume as the cube. he does the calculation below. this is really clever.

That makes sense now.

good image

this graph helps a lot .. i wasn't picturing like this

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore 1/2 *no matter what* b *and* h *are*, so $A = bh/2$.

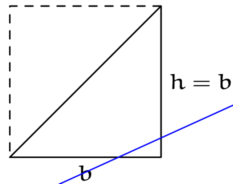

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

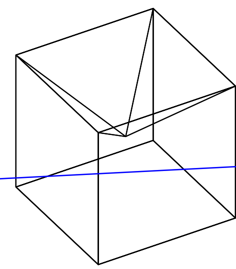

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

so where is the loss of information? i'm confused as to how this fits in this section.

I don't think there was necessarily a loss of information in this example, but it definitely showed how we took a seemingly-complex problem and solved it by simplifying it into an easy case (in this particular example, assuming that the base is twice the height and also assuming that the pyramid is square). I guess the "loss of information" would be that we assumed that the base was square and also that 2h = b. This example fits this section quite nicely and does a good job of setting the tone for the chapter.

I would look at it more as losing complexity (sometimes at the cost of losing information). in this example, we're discarding complexity by considering this "simple case" with pyramids with b=2, h=1 (the loss of information would be the limitation of this solution to this case). then, if we find a correct answer, it has to work for other cases too, so we go back to the general case with base b and height h.

How did you determine that only the ratio matters? Or was that something you just made to be so?

well i suppose h/b is a dimensionless group...

We're just trying to find a set of pyramids that will fit inside a cube. Obviously if we scale everything by some constant factor, they'll still fit inside a (scaled) cube, so the absolute size doesn't matter. Thus, only the ratio is important.

I dont really understand how we determine which aspect ratio is "right."

I guess the "right" aspect ration is the one that makes the problem the easiest. The determination of what is easiest, I suspect, is the hard part.

Aha, but you had to know h would be 1/2b to pick b=2. The process can be iterative to make the numbers easier, I know.

It really doesn't matter what you choose, though. Personally I would choose b=1, because that makes the volume trivial to compute.

Shouldn't it just be b/2? so they can meet at the point.

confusing visual for me

I agree. I had to draw a 2D picture from different viewpoints to understand.

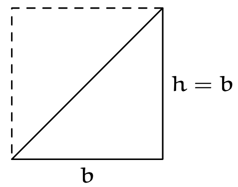

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore 1/2 *no matter what* b *and* h *are*, so $A = bh/2$.

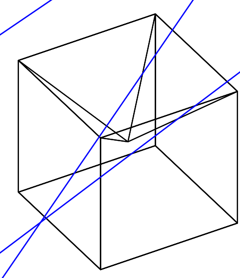

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

Very clever.

an somewhat unrelated note, i think kids in elementary schools would find this method very amusing, fun and easy to understand!!!

this is really cool but i also think that it could be explained more concisely – it's pretty intuitive with the right picture

Cool! The picture on the side helps get it across as well.

I agree, it's a nice clarification.

I agree. I got a bit lost in the description, but the picture really cleared things up.

I thought the description made more sense than the picture, so there's something for everyone to understand.

Both helped me, I think

I never realized this! Seems to simple now

Yeah, this was a great example. I always did the 2D version in my head to get the 1/2, but never extended it to get the 3D!

This explanation was a little confusing to me. I would change V to V_pyramid and I would reward the phrase "and must make volume 4/3" to "and must make V_pryamid equal to 4/3."

How did we jump to this?

For a 3-D solid, the volume will be proportional to the product of the three length dimensions, in this case of a pyramid with a square base, h*b*b = hbˆ2. We just don't know what factor (a fraction of 1, since it's smaller than a cube) the formula for this shape has.

Right, the added length from the 2D model is another base, since we are assuming the base is square.

I really like this explanation. It was very easy to follow.

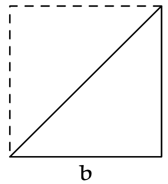

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore 1/2 *no matter what* b *and* h *are*, so $A = bh/2$.

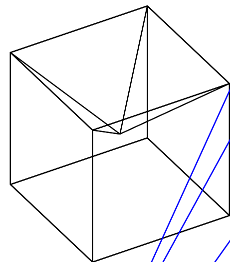

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

This is an interesting way to spatially break down the problem. I always found it easier to use 3 triangles with a 90deg corner in the space of a cube, such that all of the square bases lie on the xy, xz, and yz planes.

I'm confused about how the fact that these particular values for h and b are relevant for any pyramid.

He stated earlier that he chose b=2 for this arithmetic but in order to have six pyramids that fit into that cube, the constant must be 1/3.

The way this was arrived at was pretty easy to follow, but I think one more statement about how this easier case allowed us to find the constant for ALL cases would be useful

I'm confused as to how you can claim this is hbˆ2. How do we eliminate bˆ2/h as a possibility? We may know what the area of a pyramid should be, but in a problem where we don't know the proper solution, how do we eliminate alternatives?

That's a good point I didn't think about that alternative I'm curious as to how we can discard wrong solutions as well

I think we can intuit that by looking at the dimensions. First clue is that volume needs to have dimensions of Lˆ3. Since b and h are both L dimensions you then have to figure out whats squared and what not. Since height square doesn't really make sense for how we define height, squaring the base dimension works.

In this case, I will agree. It's always nice to summarize after a long problem.

I agree, this conclusion seems hastily drawn...a slower and longer explanation would make things perfectly clear.

I also agree that a summary might help...the example was great, however. As for the uncertainty about why we twiddle hbˆ2, look at the note surrounding the equation. It comes from the 3 dimensions the pyramid is comprised of. The same base and height as a triangle, but an added 3d dimension outward of b, since we are assuming the base is square.

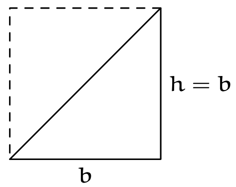

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore 1/2 *no matter what* b *and* h *are*, so $A = bh/2$.

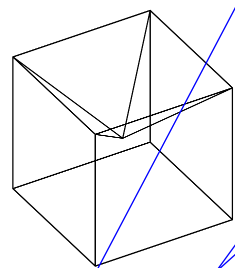

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume 4/3. Since $hb^2 = 4$ for these pyramids, the missing constant is 1/3. Voilà:

$$V = \frac{1}{3}hb^2.$$

But this would only work for pyramids where the base is twice the height

ie: An easy case.

Would it then be beneficial / interesting to show how the easy case can be extrapolated so that the base does not have to be twice the height? At least I'd like to see it.

I think the point of this section is to just prove easy cases and say "well, it's a good enough approximation for all other cases" without worrying about proving said other cases.

is this idea perhaps to be used in junction with dimensional analysis at all? or is it a separate thing?

I like this, but for a lot of our uses, won't we end up still leaving the constant out? So it would be V hb^2

Whether we are leaving the constant out in future problems (taking advantage of the dimensional analysis method), we wanted to actually solve for the volume of a square pyramid, so while we did find out that V hb^2, we were really interested in simplifying this case in order to determine that the factor is in fact 1/3. That was the goal of this example problem.

True, we might do so as a way to quickly achieve other goals, but here, the exact formula is our goal.

The point here is not to estimate a volume of some pyramid, the goal is to find an estimate for the formula.

I think in stead of just proving the case, using simple numbers to show how it works helps a lot.

I agree, though I must confess that I do constantly fight the urge to ask "But how does that show that work in EVERY case?" Easy cases seems to go against the urge to make everything like proof.

Like everyone else has said, I think this is clear and useful. It might be tempting to use this all of the time, even when it's not the best way to do a problem.

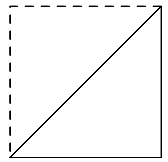

Rather than explaining right away the one-third in the volume of a pyramid, a difficult three-dimensional problem, let's first find the corresponding constant in a two-dimensional problem. That problem is the area of a triangle with base b and height h; its area is $A \sim bh$. What is the constant? Choose a convenient triangle – a special, easy case – perhaps a 45-degree right triangle where $h = b$. Two such triangles form a square with area b^2, so $A = b^2/2$ when $h = b$. The constant in $A \sim bh$ is therefore $1/2$ *no matter what* b *and* h *are*, so $A = bh/2$.

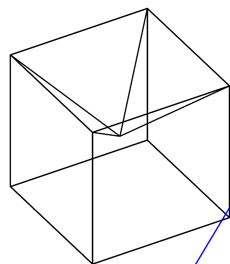

Now use the same construction in three dimensions. What square-based pyramid, when combined with itself perhaps several times, makes a familiar shape? In other words, find an easy case of h and b, and make sure that the volume is correct in that case. Because only the aspect ratio h/b matters in the following discussion, choose b conveniently, then choose h to make a pyramid with the right aspect ratio. The goal shape is suggested by the square pyramid base. One easy solid with a square base is a cube.

Therefore, let's try to combine several pyramids into a cube of side b. To simplify the upcoming arithmetic, I choose $b = 2$. What should the height h be? To decide, imagine how the cube will be constructed. Each cube has six faces, so six pyramids might make a cube where each pyramid base forms one face of the cube, and each pyramid tip faces inward, meeting in the center of the cube. For the tips to meet in the center of the cube, the height must be $h = 1$. So six pyramids with $b = 2$, and $h = 1$ make a cube with side length 2.

The volume of one pyramid is one-sixth of the volume of this cube:

$$V = \frac{\text{cube volume}}{6} = \frac{8}{6} = \frac{4}{3}.$$

The volume of the pyramid is $V \sim hb^2$, and the missing constant must make volume $4/3$. Since $hb^2 = 4$ for these pyramids, the missing constant is $1/3$. Voilà:

$$V = \frac{1}{3}hb^2.$$

I agree with what was earlier that it would be nice to see the link on how we can use easy case to understand the concpet and use it to solve more difficult cases. Unless we ware just saying do easy cases and assume them to be close enough. If that's the case then I think maybe stating something like that a little more explicitly would be helpful, that way people are not trying to make the connection.

I really like examples like this where you take basic examples we know and 1.) bring up ideas/questions we never thought about and 2.) make us realize how complex the problem can be and 3.) as always, simplify.

Yeah, I like being shown something I already know so I'm 100% sure the method works.

I agree as well. This might have been one of the cleanest sections I've read in a long while

yep it was a great example!

6.2 Atwood machine

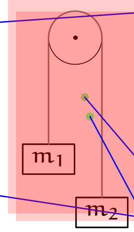

The next problem illustrates dimensional analysis and easy cases in a physical problem. The problem is the Atwood machine, a staple of the first-year physics curriculum. Two masses, m_1 and m_2, are connected and, thanks to a pulley, are free to move up and down. What is the acceleration of the masses and the tension in the string? You can solve this problem with standard methods from first-year physics, which means that you can can check the solution that we derive using dimensional analysis, easy cases, and a feel for functions.

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is to decide what variables are relevant. The acceleration depends on gravity, so g should be on the list. The masses affect the acceleration, so m_1 and m_2 are on the list. And that's it. You might wonder what happened to the tension: Doesn't it affect the acceleration? It does, but it is itself a consequence of m_1, m_2, and g. So adding tension to the list does not add information; it would instead make the dimensional analysis difficult.

Var	Dim	What
a	LT^{-2}	accel. of m_1
g	LT^{-2}	gravity
m_1	M	block mass
m_2	M	block mass

These variables fall into two pairs where the variables in each pair have the same dimensions. So there are two dimensionless groups here ripe for picking: $G_1 = m_1/m_2$ and $G_2 = a/g$. You can make any dimensionless group using these two obvious groups, as experimentation will convince you. Then, following the usual pattern,

$$\frac{a}{g} = f\left(\frac{m_1}{m_2}\right),$$

where f is a dimensionless function.

Pause a moment. The more thinking that you do to choose a clean representation, the less algebra you do later. So rather than find f using m_1/m_2 as the dimensionless group, first choose a better group. The ratio m_1/m_2 does not respect the symmetry of the problem in that only the sign of the acceleration changes when you interchange the labels m_1 and m_2. Whereas m_1/m_2 turns into its reciprocal. So the function f will have

I'm a little confused as to what exactly is the "easy case" in this problem (vs the complex case). In the pyramid example we expanded the triangle case into the harder pyramid case, but I can't seem to distinguish this step in the Atwood example.

The easy cases are taking m1 to be huge and m1=m2. There are lots of other mass ratios that we could use, but those would be harder cases.

The easy cases come at the end when he tries to check to see if the answer is correct.

Don't remember if they ever actually called it this..

My physics teacher did.

I'm glad a figure is here

I am not very comfortable with physics but the way this is explained is very helpful and clear to me.

Ew, 8.01.

Physics is a very valid place to use happy approximation techniques, otherwise it gets VERY nasty very fast.

Agreed. It's definitely easier to work with physics when combined with approximation.

If we model the cow as a sphere...

How did we come to this conclusion? Is this the second use of easy cases that we will describe in this section?

typo

unless you dance when you check the solution.

hehe... you made me laugh :)

NB should have a "Like" option in addition to the "Agree"

6.2 Atwood machine

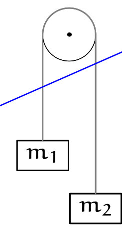

The next problem illustrates dimensional analysis and easy cases in a physical problem. The problem is the Atwood machine, a staple of the first-year physics curriculum. Two masses, m_1 and m_2, are connected and, thanks to a pulley, are free to move up and down. What is the acceleration of the masses and the tension in the string? You can solve this problem with standard methods from first-year physics, which means that you can can check the solution that we derive using dimensional analysis, easy cases, and a feel for functions.

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is to decide what variables are relevant. The acceleration depends on gravity, so g should be on the list. The masses affect the acceleration, so m_1 and m_2 are on the list. And that's it. You might wonder what happened to the tension: Doesn't it affect the acceleration? It does, but it is itself a consequence of m_1, m_2, and g. So adding tension to the list does not add information; it would instead make the dimensional analysis difficult.

Var	Dim	What
a	LT^{-2}	accel. of m_1
g	LT^{-2}	gravity
m_1	M	block mass
m_2	M	block mass

These variables fall into two pairs where the variables in each pair have the same dimensions. So there are two dimensionless groups here ripe for picking: $G_1 = m_1/m_2$ and $G_2 = a/g$. You can make any dimensionless group using these two obvious groups, as experimentation will convince you. Then, following the usual pattern,

$$\frac{a}{g} = f\left(\frac{m_1}{m_2}\right),$$

where f is a dimensionless function.

Pause a moment. The more thinking that you do to choose a clean representation, the less algebra you do later. So rather than find f using m_1/m_2 as the dimensionless group, first choose a better group. The ratio m_1/m_2 does not respect the symmetry of the problem in that only the sign of the acceleration changes when you interchange the labels m_1 and m_2. Whereas m_1/m_2 turns into its reciprocal. So the function f will have

what is the point of going through a dimensional analysis and a more standard 8.01 approach? Seems like we aren't saving anything here.

I think the point is to show how you can use it with dimensional analysis and easy cases, but then check the answer with standard 8.01 physics. You may not always be able to check your answer so easily...

And, it's still pretty easy to make a mistake when doing things by normal 8.01 procedures. So having two separate methods allows us to be fairly certain that our answer is correct.

Perhaps as a proof that the dimensional analysis method works?

Shouldn't it still be g?

I believe it should be dependent on m2 because of the rope and pulley. It is not in free fall.
If you drew a free body diagram, there would be a downward force of m1g, but there is also an upward force from the rope. You can imagine if the masses were equal, a would = 0, not g.

This notation is confusing.

I agree. You should say "equal to the acceleration of m1 or -m1" not the equal to the mass itself.

It's just saying that the sign of m_1 is the opposite of the sign of m_2.

i don't think that's what they're objecting too. more like it should be written a1 or -a1.

Agreed with the person above.

I understand that it's important to continually emphasize previously learned points, but we've just finished the unit on dimensional analysis, and personally I find it hard to wrap my head around this new topic when we're still trying to use the old dimensional analysis problem solving techniques. I'm not sure I could say what "easy cases" is if I can't separate it from dimensional analysis.

For me it helps me to get a better grasp of dimensional analysis. I think it's valuable to see it in following sections because it solidifies my understanding of it and I don't just stop using it as soon as I finish the unit.

6.2 Atwood machine

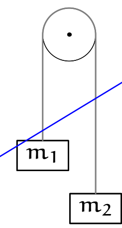

The next problem illustrates dimensional analysis and easy cases in a physical problem. The problem is the Atwood machine, a staple of the first-year physics curriculum. Two masses, m_1 and m_2, are connected and, thanks to a pulley, are free to move up and down. What is the acceleration of the masses and the tension in the string? You can solve this problem with standard methods from first-year physics, which means that you can can check the solution that we derive using dimensional analysis, easy cases, and a feel for functions.

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is to decide what variables are relevant. The acceleration depends on gravity, so g should be on the list. The masses affect the acceleration, so m_1 and m_2 are on the list. And that's it. You might wonder what happened to the tension: Doesn't it affect the acceleration? It does, but it is itself a consequence of m_1, m_2, and g. So adding tension to the list does not add information; it would instead make the dimensional analysis difficult.

Var	Dim	What
a	LT^{-2}	accel. of m_1
g	LT^{-2}	gravity
m_1	M	block mass
m_2	M	block mass

These variables fall into two pairs where the variables in each pair have the same dimensions. So there are two dimensionless groups here ripe for picking: $G_1 = m_1/m_2$ and $G_2 = a/g$. You can make any dimensionless group using these two obvious groups, as experimentation will convince you. Then, following the usual pattern,

$$\frac{a}{g} = f\left(\frac{m_1}{m_2}\right),$$

where f is a dimensionless function.

Pause a moment. The more thinking that you do to choose a clean representation, the less algebra you do later. So rather than find f using m_1/m_2 as the dimensionless group, first choose a better group. The ratio m_1/m_2 does not respect the symmetry of the problem in that only the sign of the acceleration changes when you interchange the labels m_1 and m_2. Whereas m_1/m_2 turns into its reciprocal. So the function f will have

so we cant always discard mass in gravity problems. you said that in class the other day and it had been bugging me.

It doesn't matter in the general case of accelerating towards earth, as Galileo showed in his experiment dropping two objects. It does matter in this case, however, since the two objects are applying a force on each other, not an acceleration. Think about it...what if the two objects were the same mass? What if 1 object was weightless?

I should be more precise. You can discard the mass – if there's only one mass. With only one mass, it cannot be part of a dimensionless group.

But with two masses, their ratio is dimensionless, so the masses can affect the behavior (but only through their ratio, not their actual values – i.e. if you double each mass, nothing changes).

That's were it gets tricky though – knowing that's it. I feel like if I were to do this I'd think I'd forgotten something...

(to clarify, I wasn't thinking of tension so much as friction of the pulley, etc. – are we just neglecting those?)

I think we are neglecting friction, but I agree that it's hard to know exactly where to stop. This almost certainly qualifies as a 'use the force' moment, which seems to just be something you have to get a sense for.

I think this 'that's it' decision comes from a feel of what we want to get out of the problem. We only put in what we want to get out. We don't care so much about how friction might affect the solution.

I never would have thought of tension- it's pretty obvious that it's not an external force and would therefore be dependent on m and g

right. it's sort of like if you drew a control volume around the system, it's not an external force acting on the system.

6.2 Atwood machine

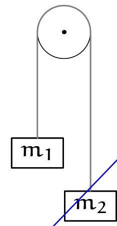

The next problem illustrates dimensional analysis and easy cases in a physical problem. The problem is the Atwood machine, a staple of the first-year physics curriculum. Two masses, m_1 and m_2, are connected and, thanks to a pulley, are free to move up and down. What is the acceleration of the masses and the tension in the string? You can solve this problem with standard methods from first-year physics, which means that you can can check the solution that we derive using dimensional analysis, easy cases, and a feel for functions.

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is to decide what variables are relevant. The acceleration depends on gravity, so g should be on the list. The masses affect the acceleration, so m_1 and m_2 are on the list. And that's it. You might wonder what happened to the tension: Doesn't it affect the acceleration? It does, but it is itself a consequence of m_1, m_2, and g. So adding tension to the list does not add information; it would instead make the dimensional analysis difficult.

Var	Dim	What
a	LT^{-2}	accel. of m_1
g	LT^{-2}	gravity
m_1	M	block mass
m_2	M	block mass

These variables fall into two pairs where the variables in each pair have the same dimensions. So there are two dimensionless groups here ripe for picking: $G_1 = m_1/m_2$ and $G_2 = a/g$. You can make any dimensionless group using these two obvious groups, as experimentation will convince you. Then, following the usual pattern,

$$\frac{a}{g} = f\left(\frac{m_1}{m_2}\right),$$

where f is a dimensionless function.

Pause a moment. The more thinking that you do to choose a clean representation, the less algebra you do later. So rather than find f using m_1/m_2 as the dimensionless group, first choose a better group. The ratio m_1/m_2 does not respect the symmetry of the problem in that only the sign of the acceleration changes when you interchange the labels m_1 and m_2. Whereas m_1/m_2 turns into its reciprocal. So the function f will have

I think throwing in the buzz word that tension is "dependent" on the other variables. Describing independence helped clarify in the last reading.

Does this even work when we take friction into account?

This assumes a frictionless cord/pulley.

I found it easy to see how tension doesn't matter when I realized that the tension on the m1 side of the pulley canceled out the tension on the m2 side.

I don't know if it's my screen but this line appears to be in a lighter shade of gray than the black

I think it's just your screen

mine too.

It was once in red (before I made the book all black and white), to indicate that "a" is the quantity we are looking for. I think the red got grayscaled, meaning that the color turned into a shade of gray. I'll find another way to indicate that "a" is the goal variable.

I think boldface might do the trick. Or worst case just leave it plain... I don't think it necessarily needs to be called out. (Looking back, it appears the same problem occurred in the tables of other readings, too.)

I really like that we have this box here. It makes thinking about the problem much easier.

I agree, it would be nice to see boxes like this in other parts of the readings.

I agree. These boxes make it really really easy to see the relationships between variables.

How do we know it's a/g and not g/a?

We don't. g/a would be the same group, to the -1 power. As he mentions in the next sentence, we can make any dimensionless groups from these two groups.

I didn't realize this meant "group 1" right away, I was looking for a variable G or universal gravitational constant.

Haha. That's what i did too.

6.2 Atwood machine

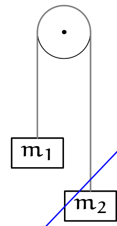

The next problem illustrates dimensional analysis and easy cases in a physical problem. The problem is the Atwood machine, a staple of the first-year physics curriculum. Two masses, m_1 and m_2, are connected and, thanks to a pulley, are free to move up and down. What is the acceleration of the masses and the tension in the string? You can solve this problem with standard methods from first-year physics, which means that you can can check the solution that we derive using dimensional analysis, easy cases, and a feel for functions.

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is to decide what variables are relevant. The acceleration depends on gravity, so g should be on the list. The masses affect the acceleration, so m_1 and m_2 are on the list. And that's it. You might wonder what happened to the tension: Doesn't it affect the acceleration? It does, but it is itself a consequence of m_1, m_2, and g. So adding tension to the list does not add information; it would instead make the dimensional analysis difficult.

These variables fall into two pairs where the variables in each pair have the same dimensions. So there are two dimensionless groups here ripe for picking: $G_1 = m_1/m_2$ and $G_2 = a/g$. You can make any dimensionless group using these two obvious groups, as experimentation will convince you. Then, following the usual pattern,

Var	Dim	What
a	LT^{-2}	accel. of m_1
g	LT^{-2}	gravity
m_1	M	block mass
m_2	M	block mass

$$\frac{a}{g} = f\left(\frac{m_1}{m_2}\right),$$

where f is a dimensionless function.

Pause a moment. The more thinking that you do to choose a clean representation, the less algebra you do later. So rather than find f using m_1/m_2 as the dimensionless group, first choose a better group. The ratio m_1/m_2 does not respect the symmetry of the problem in that only the sign of the acceleration changes when you interchange the labels m_1 and m_2. Whereas m_1/m_2 turns into its reciprocal. So the function f will have

is this equivalent to writing a/g (m_1/m_2)?

Not necessarily, since we don't know that it's a linear function. It could be a/g = C*(m1/m2)^2

It's the same as discussed in the last reading memo, where if you have 2 dimensionless groups, 1 is a function of the other, it doesn't specify how it is a function

No. The two might not be linearly related (the twiddle just hides a constant). The function could be linear, or quadratic, or inverse, or some other power of the second group

I'm still unsure how this differs from dimensional analysis.

you could also maybe write or G2 = f(G1) just to remind us of this replacement notation that is later used

This is a much clearer example of a dimensionless function than the previous readings.

I agree. The choice of dimensionless groups in the above paragraph is made very clear as well

What would a function with dimensions be?

I like this idea.

Yeah, this is a very helpful principle in general.

It's analogous to, "The more triage that you do before you pack for a trip, the easier your whole trip will be."

I figured this out the hard way many times.

Same here! I remember erasing profusely on these questions in 8.01

How can we see that this group isn't already good enough?

This is tricky. Looking at the symmetry helps, but it's not something I would have thought of at first.

6.2 Atwood machine

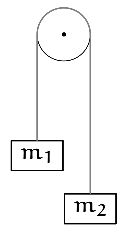

The next problem illustrates dimensional analysis and easy cases in a physical problem. The problem is the Atwood machine, a staple of the first-year physics curriculum. Two masses, m_1 and m_2, are connected and, thanks to a pulley, are free to move up and down. What is the acceleration of the masses and the tension in the string? You can solve this problem with standard methods from first-year physics, which means that you can can check the solution that we derive using dimensional analysis, easy cases, and a feel for functions.

The first problem is to find the acceleration of, say, m_1. Since m_1 and m_2 are connected by a rope, the acceleration of m_2 is, depending on your sign convention, either equal to m_1 or equal to $-m_1$. Let's call the acceleration a and use dimensional analysis to guess its form. The first step is to decide what variables are relevant. The acceleration depends on gravity, so g should be on the list. The masses affect the acceleration, so m_1 and m_2 are on the list. And that's it. You might wonder what happened to the tension: Doesn't it affect the acceleration? It does, but it is itself a consequence of m_1, m_2, and g. So adding tension to the list does not add information; it would instead make the dimensional analysis difficult.

Var	Dim	What
a	LT^{-2}	accel. of m_1
g	LT^{-2}	gravity
m_1	M	block mass
m_2	M	block mass

These variables fall into two pairs where the variables in each pair have the same dimensions. So there are two dimensionless groups here ripe for picking: $G_1 = m_1/m_2$ and $G_2 = a/g$. You can make any dimensionless group using these two obvious groups, as experimentation will convince you. Then, following the usual pattern,

$$\frac{a}{g} = f\left(\frac{m_1}{m_2}\right),$$

where f is a dimensionless function.

Pause a moment. The more thinking that you do to choose a clean representation, the less algebra you do later. So rather than find f using m_1/m_2 as the dimensionless group, first choose a better group. The ratio m_1/m_2 does not respect the symmetry of the problem in that only the sign of the acceleration changes when you interchange the labels m_1 and m_2. Whereas m_1/m_2 turns into its reciprocal. So the function f will have

Easy cases, dimensional analysis, and now symmetry? Looks like we're using many tools to solve this problem!

It's like a meta-problem! I do like that he's back referencing previous units though. It makes the order we learned them in make more sense

Yep, that seems to be the case. Admittedly, I stlil probably would ot approach this correctly though.

I'm just confused as to what symmetry exactly we're violating.

I agree, I feel like we're using too many tools. Sometimes I like references to previous units, but this time it feels like too much. The references to the previous units are overpowering the relevant "easy cases" content.

The symmetry in the way the problem is set up states that if a/g=f(m1/m2), then f(1/x)=-f(x). (i.e. switching the masses reverses the acceleration.) This is a much harder function to construct than one in which, say, f(-x)=-f(x), as in the bending light example.

i'm confused about what this means...

how do you know this? would it still work? are we looking for a new dimensionless group or just going to re arrange the ones we have?

this sentence seems awkward by itself; maybe combine it with the last sentence via a comma?

I don't understand what he's trying to say.

Agreed, it doesn't really seem to connect well with the previous section.

Agreed it could be worded better, but in case you didn't catch it from the rest of the reading: He's saying that the way the function is written doesn't match the way the system acts in real life. In real life if you switched the masses attached to the pulley, it would accelerate in the opposite direction. The way the current function is written, switching the masses creates a reciprocal instead of switching the sign of the acceleration.

i definitely did not get that from this sentence.

to do lots of work to turn the unsymmetric ratio m_1/m_2 into a symmetric acceleration.

Back to the drawing board for how to fix G_1. Another option is to use $m_1 - m_2$. Wait, the difference is not dimensionless! I fix that problem in a moment. For now observe the virtue of $m_1 - m_2$. It shows a physically reasonable symmetry under mass interchange: $G_1 \to -G_1$. To make it dimensionless, divide it by another mass. One candidate is m_1:

$$G_1 = \frac{m_1 + m_2}{m_1}.$$

That choice, like dividing by m_2, abandons the beloved symmetry. But dividing by $m_1 + m_2$ solves all the problems:

$$G_1 = \frac{m_1 - m_2}{m_1 + m_2}.$$

This group is dimensionless and it respects the symmetry of the problem.

Using this G_1, the solution becomes

$$\frac{a}{g} = f\left(\frac{m_1 - m_2}{m_1 + m_2}\right),$$

where f is another dimensionless function.

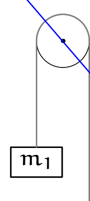

To guess $f(x)$, where $x = G_1$, try the easy cases. First imagine that m_1 becomes huge. A quantity with mass cannot be huge on its own, however. Here huge means *huge relative to* m_2, whereupon $x \approx 1$. In this thought experiment, m_1 falls as if there were no m_2 so $a = -g$. Here we've chosen a sign convention with positive acceleration being upward. If m_2 is huge relative to m_1, which means $x = -1$, then m_2 falls like a stone pulling m_1 upward with acceleration $a = g$. A third limiting case is $m_1 = m_2$ or $x = 0$, whereupon the masses are in equilibrium so $a = 0$.

Here is a plot of our knowledge of f:

That is a good point. But is it intuitive like this case most of the times? I feel like I might have trouble figuring it out when it comes to real problems.

Would this even be possible?

G1?

Why aren't you using the capital Pi from class?

referring to group 1 = m_1 / m_2

I know, but I was wondering why he didn't use the capital Pi he used in class. It seems like his variables for 'groups' are shifting. And I don't like G. It does look like the gravitational constant.

isn't this not dimensionless?

I like this process of dimensionless analysis, but I'm still uncomfortable with it sometimes. It feels like it requires a lot of creativity, or it feels like cheating...or maybe using creativity to cheat, lol.

Haha that's exactly how I feel about it...I feel like there would definitely be times where I'd neglect a variable and make some strange and wonky dimensionless groups.

I agree. I often find myself including unneeded or redundant variables.

I feel like using this method will become easier and more intuitive with practice.

It just seems strange that we can use this instead, just because it's symmetrical and it won't affect the outcome

We have to start somewhere and this is just kind of the starting point. In order to get it to be dimensionless, we will have to do more things to it later.

i don't understand what this sentence means.

I don't either; why exactly do you choose m1 - m2, why not m2 - m1, or m1+m2? It is just because m1 is moving in the opposite direction of m2? Then, why do you divide by m1+m2 lower down?

to do lots of work to turn the unsymmetric ratio m_1/m_2 into a symmetric acceleration.

Back to the drawing board for how to fix G_1. Another option is to use $m_1 - m_2$. Wait, the difference is not dimensionless! I fix that problem in a moment. For now observe the virtue of $m_1 - m_2$. It shows a physically reasonable symmetry under mass interchange: $G_1 \to -G_1$. To make it dimensionless, divide it by another mass. One candidate is m_1:

$$G_1 = \frac{m_1 + m_2}{m_1}.$$

That choice, like dividing by m_2, abandons the beloved symmetry. But dividing by $m_1 + m_2$ solves all the problems:

$$G_1 = \frac{m_1 - m_2}{m_1 + m_2}.$$

This group is dimensionless and it respects the symmetry of the problem.

Using this G_1, the solution becomes

$$\frac{a}{g} = f\left(\frac{m_1 - m_2}{m_1 + m_2}\right),$$

where f is another dimensionless function.

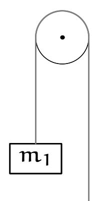

To guess $f(x)$, where $x = G_1$, try the easy cases. First imagine that m_1 becomes huge. A quantity with mass cannot be huge on its own, however. Here huge means *huge relative to* m_2, whereupon $x \approx 1$. In this thought experiment, m_1 falls as if there were no m_2 so $a = -g$. Here we've chosen a sign convention with positive acceleration being upward. If m_2 is huge relative to m_1, which means $x = -1$, then m_2 falls like a stone pulling m_1 upward with acceleration $a = g$. A third limiting case is $m_1 = m_2$ or $x = 0$, whereupon the masses are in equilibrium so $a = 0$.

Here is a plot of our knowledge of f:

I have no idea what you mean by this little figure...

I think its just to show that acceleration symmetry mentioned earlier on, since acceleration depends on m1 and m2

If the values of m1 and m2 are switched, then the acceleration of m1 will be of the same magnitude, but in the opposite direction, so the sign will change. That's what "mass interchange" means, though it's a poor term.

This is a cool way to thing about the physics without really thinking about the physics.

I remembered the solutions in 8.01 always had this sort of form!

so in the examples in the previous readings/in class, we could have done this too, like (var 1+ var2)/var1 or other types of manipulations such that it ends up to be dimensionless, why didn't we do all these before, why did we just use simple multiplication or division?

to do lots of work to turn the unsymmetric ratio m_1/m_2 into a symmetric acceleration.

Back to the drawing board for how to fix G_1. Another option is to use $m_1 - m_2$. Wait, the difference is not dimensionless! I fix that problem in a moment. For now observe the virtue of $m_1 - m_2$. It shows a physically reasonable symmetry under mass interchange: $G_1 \to -G_1$. To make it dimensionless, divide it by another mass. One candidate is m_1:

$$G_1 = \frac{m_1 + m_2}{m_1}.$$

That choice, like dividing by m_2, abandons the beloved symmetry. But dividing by $m_1 + m_2$ solves all the problems:

$$G_1 = \frac{m_1 - m_2}{m_1 + m_2}.$$

This group is dimensionless and it respects the symmetry of the problem.

Using this G_1, the solution becomes

$$\frac{a}{g} = f\left(\frac{m_1 - m_2}{m_1 + m_2}\right),$$

where f is another dimensionless function.

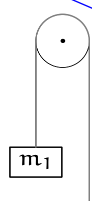

To guess $f(x)$, where $x = G_1$, try the easy cases. First imagine that m_1 becomes huge. A quantity with mass cannot be huge on its own, however. Here huge means *huge relative to* m_2, whereupon $x \approx 1$. In this thought experiment, m_1 falls as if there were no m_2 so $a = -g$. Here we've chosen a sign convention with positive acceleration being upward. If m_2 is huge relative to m_1, which means $x = -1$, then m_2 falls like a stone pulling m_1 upward with acceleration $a = g$. A third limiting case is $m_1 = m_2$ or $x = 0$, whereupon the masses are in equilibrium so $a = 0$.

Here is a plot of our knowledge of f:

so in the examples in the previous readings/in class, we could have done this too, like (var 1+ var2)/var1 or other types of manipulations such that it ends up to be dimensionless, why didn't we do all these before, why did we just use simple multiplication or division?

yeah if this was always allowed, then I feel like we should have done an example like this in which the variables are added or subtracted in the unit about dimensional analysis.

I think this problem is harder then some of the other ones we have done so we need to use a bit more complex method. We cannot just divide or multiply so it becomes dimensionless.

I agree, but in this case there are clearly two separate variables of the same dimensionality, so it makes more sense. Before, you would have had to make strange groups like (2F - ma)/F, which would just give a dimensionless constant. It usually makes more sense to multiply to get dimensionless groups unless you have two separate things like lift and drag or two different masses.

The rough intuition is that the more closely the dimensionless group matches the physics of the problem, the less work the dimensionless function f has to do. In other words, the simpler the function f is, and the easier it is to guess.

So, the whole approach is a kind of divide and conquer. Instead of using the simplest group, e.g. m1/m2, and then having a hard problem to guess f, you do some work to turn m1/m2 into a nicer group, and then some work to guess f.

This is all fine to understand once it's already done, but I sincerely doubt I would go through the same rationale to get the result that ends up working.

I definitely feel that way the first time I see a problem solved a new way. But I think that once you see it done a couple of times, you figure out how to apply it to solving your own problems.

hopefully he'll give a few other examples in class, and this will become more intuitive

and now we know to keep it in mind as we solve other problems.

True - sometimes on the problem sets I see how a concept could be applied, but I often feel myself struggling because I've missed some important step.

to do lots of work to turn the unsymmetric ratio m_1/m_2 into a symmetric acceleration.

Back to the drawing board for how to fix G_1. Another option is to use $m_1 - m_2$. Wait, the difference is not dimensionless! I fix that problem in a moment. For now observe the virtue of $m_1 - m_2$. It shows a physically reasonable symmetry under mass interchange: $G_1 \to -G_1$. To make it dimensionless, divide it by another mass. One candidate is m_1:

$$G_1 = \frac{m_1 + m_2}{m_1}.$$

That choice, like dividing by m_2, abandons the beloved symmetry. But dividing by $m_1 + m_2$ solves all the problems:

$$G_1 = \frac{m_1 - m_2}{m_1 + m_2}.$$

This group is dimensionless and it respects the symmetry of the problem.

Using this G_1, the solution becomes

$$\frac{a}{g} = f\left(\frac{m_1 - m_2}{m_1 + m_2}\right),$$

where f is another dimensionless function.

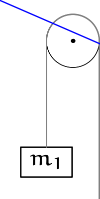

To guess $f(x)$, where $x = G_1$, try the easy cases. First imagine that m_1 becomes huge. A quantity with mass cannot be huge on its own, however. Here huge means *huge relative to* m_2, whereupon $x \approx 1$. In this thought experiment, m_1 falls as if there were no m_2 so $a = -g$. Here we've chosen a sign convention with positive acceleration being upward. If m_2 is huge relative to m_1, which means $x = -1$, then m_2 falls like a stone pulling m_1 upward with acceleration $a = g$. A third limiting case is $m_1 = m_2$ or $x = 0$, whereupon the masses are in equilibrium so $a = 0$.

Here is a plot of our knowledge of f:

It almost feels like "easy cases" is taking a back seat to dimensional analysis here. I realize you need to do dimensional analysis to get to the easy cases part, but is there a way to condense the DA to focus on the easy cases?

I agree - this seems like we're just doing work from the previous section rather than working on "easy cases"

I dunno, I kind of like this, as it's illustrating how the previous topic interplays with the current topic. We don't know much about "easy cases", so it makes sense that we'll need to rely on our other techniques at first. Plus, this actually tied together a few more knots for me with regards to dimensional analysis

I am a little confused on what you mean by symmetry here

Why does putting m_1 + m_2 keep symmetry?

does this violate the equation? dividing one side but not the other?

Not really – remember, we're still using an unknown function f() to relate one side to the other. This just means that f() will be different.

No, we are still working inside the f(). As long as that remains dimensionless, the equation is satisfied.

Interchanging m1 and m2 should give an acceleration with the opposite sign because of the way the system is physically connected. In order to make our function f() as simple as possible, we are making the argument to f() similarly symmetric.

I forgot to take momentum into account when I first thought about this.

how does it respect he symmetry of the problem, can you explain this a bit more

it respects the symmetry, in that if you interchange m1 with m2, you get the negative of the exact same value (which we want–intuitively, this makes sense).

i really like this idea of "the symmetry of the problem." AKA we know that nature doesn't care about our imposed coordinates systems or signs.

@ previous post: that was a really interested way of looking at it, that nature doesn't care about our imposed coordinate systems or signs, and the symmetry of problems. That puts it in a great perspective!

to do lots of work to turn the unsymmetric ratio m_1/m_2 into a symmetric acceleration.

Back to the drawing board for how to fix G_1. Another option is to use $m_1 - m_2$. Wait, the difference is not dimensionless! I fix that problem in a moment. For now observe the virtue of $m_1 - m_2$. It shows a physically reasonable symmetry under mass interchange: $G_1 \to -G_1$. To make it dimensionless, divide it by another mass. One candidate is m_1:

$$G_1 = \frac{m_1 + m_2}{m_1}.$$

That choice, like dividing by m_2, abandons the beloved symmetry. But dividing by $m_1 + m_2$ solves all the problems:

$$G_1 = \frac{m_1 - m_2}{m_1 + m_2}.$$

This group is dimensionless and it respects the symmetry of the problem.

Using this G_1, the solution becomes

$$\frac{a}{g} = f\left(\frac{m_1 - m_2}{m_1 + m_2}\right),$$

where f is another dimensionless function.

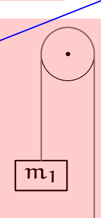

To guess $f(x)$, where $x = G_1$, try the easy cases. First imagine that m_1 becomes huge. A quantity with mass cannot be huge on its own, however. Here huge means *huge relative to* m_2, whereupon $x \approx 1$. In this thought experiment, m_1 falls as if there were no m_2 so $a = -g$. Here we've chosen a sign convention with positive acceleration being upward. If m_2 is huge relative to m_1, which means $x = -1$, then m_2 falls like a stone pulling m_1 upward with acceleration $a = g$. A third limiting case is $m_1 = m_2$ or $x = 0$, whereupon the masses are in equilibrium so $a = 0$.

Here is a plot of our knowledge of f:

I am kind of confused with this - it seems like the groups are just chosen at random. How do we chose without having a lot of prior knowledge of the atwood physics.

These formulas are all the same if you were to do it algebraically, I don't see the benefit of using this new method (it's somewhat confusing)

I agree, I remember finding the derivation rather helpful in 8.01

I think the point here is to give an example of how you could use this method to solve a problem we know how to solve so that when we get tougher problems we'll know what to do. Since we already understand how to solve these problems we can trust the answers we get here and understand the reasoning better than if we were given something that we had never seen before.

This is very clearly explained. Even when I thought I had questions or problems, all I had to do was go back and re-read the particular part and my question was answered.

I don't understand why x is introduced here, it just seems to make things harder to understand.

This is very helpful because we can easily plug in m2=0, m1=0, or m1=m2, to make sure the magnitudes of our answers are relatively correct

This is always the "last problem" on any part of any physics test. "What sort of behavior will you expect if m approaches infinity? if m approaches zero? Does your equation give the expected results?"

I agree that this was very helpful. It allows us to get a tangible feeling for the physics governing the problem.

I hate to be a brat, but do you think you could avoid highlighting an entire paragraph in your comments? It makes it so we can't highlight specific portions to comment on. Thank you and happy approximating!

I too thought of limits when I read this paragraph. Isn't this one of the upcoming chapters?

to do lots of work to turn the unsymmetric ratio m_1/m_2 into a symmetric acceleration.

Back to the drawing board for how to fix G_1. Another option is to use $m_1 - m_2$. Wait, the difference is not dimensionless! I fix that problem in a moment. For now observe the virtue of $m_1 - m_2$. It shows a physically reasonable symmetry under mass interchange: $G_1 \to -G_1$. To make it dimensionless, divide it by another mass. One candidate is m_1:

$$G_1 = \frac{m_1 + m_2}{m_1}.$$

That choice, like dividing by m_2, abandons the beloved symmetry. But dividing by $m_1 + m_2$ solves all the problems:

$$G_1 = \frac{m_1 - m_2}{m_1 + m_2}.$$

This group is dimensionless and it respects the symmetry of the problem.

Using this G_1, the solution becomes

$$\frac{a}{g} = f\left(\frac{m_1 - m_2}{m_1 + m_2}\right),$$

where f is another dimensionless function.

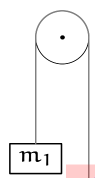

To guess $f(x)$, where $x = G_1$, try the easy cases. First imagine that m_1 becomes huge. A quantity with mass cannot be huge on its own, however. Here huge means *huge relative to* m_2, whereupon $x \approx 1$. In this thought experiment, m_1 falls as if there were no m_2 so $a = -g$. Here we've chosen a sign convention with positive acceleration being upward. If m_2 is huge relative to m_1, which means $x = -1$, then m_2 falls like a stone pulling m_1 upward with acceleration $a = g$. A third limiting case is $m_1 = m_2$ or $x = 0$, whereupon the masses are in equilibrium so $a = 0$.

Here is a plot of our knowledge of f:

It might still be nice to see an 'm2' here, especially since this paragraph also talks about a case when m1=m2 and not just the case when m1>>m2

I agree, a more detailed figure then this might be even more helpful. Possibly showing the result of the easy cases.

I think an arrow may be helpful, or some sort of force directional.

I liked how this section was a good synthesis of the things we've learned, but I feel that the easy cases part was thrown in at the end like an afterthought, rather than the main substance of this section. Though, that might be pretty representative of how easy cases are usually used.

these simple cases are excellent proof-of-concept demonstrators. will we see an example of a fluid mechanics problem?

This is not technically a plot of our knowledge of f. Such a plot would have only the 3 easy points described. Drawing a straight line connecting them is our "educated guess" as you said below. It might be nice to literally draw what you say; or you could rephrase what you said to match the plot... "say what you mean, mean what you say"

I agree...how exactly would this be a plot of our knowledge? would it be in relation to the problem that we are knowlegable about?

Also agree. There should be a step where the leap from three collinear points to the assumption a linear function occurs. Because it could be, for example, f(x)=x^k for any odd k, or f(x)=sin(2x/pi), etc.

This line would look better if it were on the next page, above the plot.

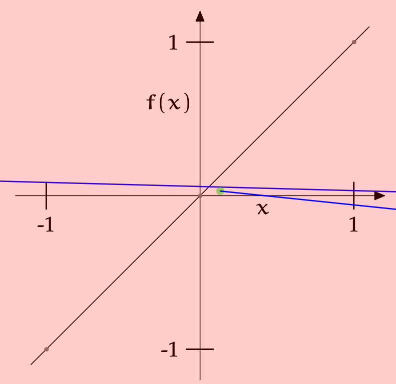

The simplest conjecture – an educated guess – is that $f(x) = x$. Then we have our result:

$$\frac{a}{g} = \frac{m_1 - m_2}{m_1 + m_2}.$$

Look how simple the result is when derived in a symmetric, dimensionless form using easy cases!

OK I see how m_1-m_2 solves this now.

Plotting the data definitely puts the info into a more readable format

Yay for diagrams!

it's be nice to plot the points we just discussed that we used to form this graph

I always found it cool how plots make things so much easier to see.

definitely - especially when they're nice linear relationships like this one

Yes, it especially helps to check the special cases and extremes on the graph and make sure they make sense to the problem.

Regarding the question above if we plot our our 3 simple cases we get three points, but how do we know that the the relationship is linear? Is it just a case of the simplest guess is usually correct?

This was just a set-up, right, for the less elegant problems we're going to encounter shortly?

Yea are all the cases this "easy?"

this clears up how we solve for f from the previous readings

I intuitively understand this, however, another step would have been nice to get from the graph to the ultimate equation.

I don't see how this is valuable enough to extend to solving most types of problems, this seems to only work when the solution is very simple anyway

Well this is how you would approach these types of problems, if you recall from 8.01. The answer / relationship here is that same that you would get fro mforce diagrams. That's really what we did, we just didn't draw them out.

I thought that too, but I believe the whole point was to show how we can use common knowledge and previous techniques to simplify problems.

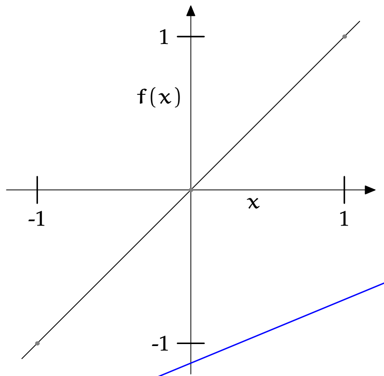

The simplest conjecture – an educated guess – is that $f(x) = x$. Then we have our result:

$$\frac{a}{g} = \frac{m_1 - m_2}{m_1 + m_2}.$$

Look how simple the result is when derived in a symmetric, dimensionless form using easy cases!

I think we need to check m1=m2 to show that the correct dimensionless quantity is a/g and not g/a?

If the dimensionless quantity was g/a, the plot would asymptotically approach -infinity in the left half plane, and +infinity in the right half plane as x tended towards zero, making it a not linear relationship.

how does this differ from the case where your second dimension less group was m/m? I don't really see the benefit of changing the groups here

This section was awesome. It's great to see how all the concepts we've learned so far can be put together and applied to solve hard problems.

I still think the algebraic method would have been easier in this case.

Well, I think the point is to show a simple case so you can then use this for harder cases. (like, rewriting G to make is symmetrical seems like a bit of extra work for a problem this simple, but it makes sense for more complicated problems)

I agree. I think the algebraic method would have been faster and easier for this one.

Actually, I disagree. I think we're comparing an algebraic method which we've been trained on for years to a new method that we've only known for a few weeks at most. Imagine how long it would have taken to go through this rigorously if we hadn't all had years of similar math and physics problems.

It looks like their is a lot of trial and error in these cases. I guess that's why it is easy cases.

so what would be the easy case here? the fact that we were choosing which variables to look at in an easy way?

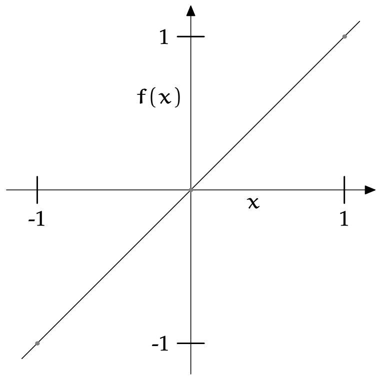

The simplest conjecture – an educated guess – is that $f(x) = x$. Then we have our result:

$$\frac{a}{g} = \frac{m_1 - m_2}{m_1 + m_2}.$$

Look how simple the result is when derived in a symmetric, dimensionless form using easy cases!

After reading this section, I'm still not sure where the "loss" is. It seems we reasoned through things, using easy cases, and we came out with pretty solid formulas.

Agreed - what was lost? I thought we were pretty thorough and looked like we came to a reasonably correct solution.

In this case I don't think anything was lost but I can imagine cases where we can not predict f(x) exactly, and maybe this is where the loss comes in?

That's true; in this case our guess was 100% correct. But it's not guaranteed to be right. And in the upcoming fluid-mechanics examples, you'll see that you get reasonably good predictions, but not exact ones – because we've thrown away information (by assuming that we are in one of the special cases).

So in simple cases there is no loss, but in regressing to simple cases from more complicated ones we lose information. ok.

But is something like throwing out friction a loss here? The two easy cases of no mass on the other end (or no rope) and equal masses would remain the same, but friction would change the shape of our plot.

This was no more lossy than the bending of light problem, for example. In all of these cases we have performed dimensional analysis and then guessed a (usually conveniently linear) function. Also, in both cases, we have predicted values for the coefficient of the linear term (although in the bending of light problem, there were several theories).

Perhaps it would serve to calm any slight ambivalence the reader might have by stating explicitly:".. in this case our guess was 100% correct. But it's not guaranteed to be right. And in the upcoming fluid-mechanics examples, you'll see that you get reasonably good predictions, but not exact ones "

I wish I had this text available when I was in physics! This concept seemed so difficult and complex when I learned it in 8.01 but after reading this section, I feel so much more confident with the concepts and feel like I actually understand them. This is one of the most concise and complete sections of any of our assignments.

I think this is interesting, but it seems like easy cases like this may as well just be done out the bruce force method.

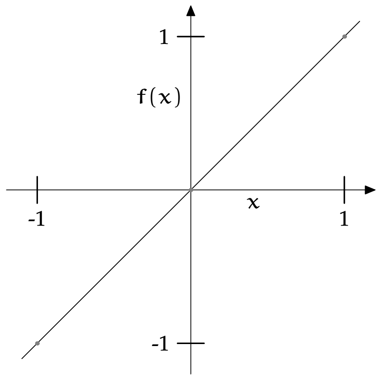

The simplest conjecture – an educated guess – is that $f(x) = x$. Then we have our result:

$$\frac{a}{g} = \frac{m_1 - m_2}{m_1 + m_2}.$$

Look how simple the result is when derived in a symmetric, dimensionless form using easy cases!

So what is "easy cases"? After reading this introduction to the topic, I'm still not sure exactly what it is...

yeah i agree..I feel like this was more about dimensionless groups. I'm not sure I understand what easy cases are.

Easy is not trying to think about cases like m1 = 40 and m2 = 60 and what the acceleration would be, but rather thinking about m1 = 40 and m2 = 0 or 40. It could also be called extreme cases here. Then we can find the solution for all the intermediate states from those two situations + symmetry.

I think doing a bit more of a formal intro to Easy Cases in the beginning of this lecture would be helpful. I realize you might reach some conclusion at the end of the section, but I think it'd be nice to have some guiding/overaching principle from the beginning.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{6.1}$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

This section is too short. I'd like to see 6.3.1 and 6.3.2 both expanded with some more explanation. I'm not lost, but just unfulfilled.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

For Sunday's memo, read this introduction to easy cases and drag (for high and low Reynolds number).

is there a reason we talk about drag so much? I feel like it comes up in every unit

what easy case that you refer to?

I believe he is referring to easy cases in general.

I also don't think there are separate categories within easy cases, so as was mentioned before, it's just easy cases in general for solving complex problems. Whenever we are using an easy, correct solution to suffice for all problems of considerable complexity, there will usually be a loss of some accuracy. That's what easy cases refers to.

Easy cases is one approach to solving difficult problems wherein you consider a simpler scenario and use the technique there to extrapolate how to approach the complex one. A basic example is considering something in 2D, and then adding another dimension to deal with a similar problem in 3D. Unfortunately, using a basic solution in a simpler case, you will sacrifice accuracy in your approach.

maybe say, "meaning, it throws..."

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

Aren't all the estimation methods we have learned somewhat lossy? How is the easy cases method so different?

nope! some cases are "lossless." for example, symmetry, like how Gauss solved the addition from 1 to 100 problem. the answer he got was EXACTLY correct and didn't loose any info at all :)

I think he means we intentionally discard something with this method. We know that we are throwing something away when we choose some sort of an extreme case. When we use methods we don't explicitly lose information, there is just some "guessing" involved.

I agree with 7:36...was there an explanation of this concept at the start of the book? either way, it might be useful to restate with you mean by "lossy method" at the start of this chapter.

Is there a better word than lossy here? I mean, I do understand it, but I winced a bit reading it. Is a method with loss, is a dropping method; I'm sure, but lossy feels weird.

lossy is a good way to describe this, although it might be good to explain it at some point to those who do not know what it means

I don't know what lossy means...

I think it means we disregard a lot of information to simplify the problem and we overlook some of the actual parts of it, just piecing it together with our own estimation

I interpreted it as unreliable. like the above said we are making many estimates along the way so things are really not that accirate.

what was the information it threw away?

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \qquad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

It might have been nice to mention this in the above examples - some people were confused before as to where the information was "lost".

I agree with this. Even though we mentioned we were moving onto "lossy" methods in class I don't think it was mentioned in the previous sections

Yup. This came up in class as well, and I also agree that it's a good point. The previous examples were lossless!

I think we were suppose to infer from the introduction for the previous section that we were moving into lossy territory.

Right, but it still seems strange to label something as such when it isn't. It's also confusing because there are so many labels of different techniques when methodologies can be similar.

I get what you are saying here, but the wording of this paragraph is kind of awkward.

I am not quite sure of the relationship between easy cases to approximation? Do we eventually build off the easy cases to solve more complex cases? or am I confused by the language?

"easy cases" is the title of this particular approximation method. Although, I feel like there has to be a better title for it...it seems to confuse a lot of people.

So much drag... why is this class considered course 6 at all?

Course 2 kids said the same thing when we were doing UNIX.

it's about the methods and I personally like using the same example as it helps me think about the problem when I know the answer as it helps me learn the method.

I agree, it's helpful to use the same problem as an example for 2 main reasons, 1: It shows that one problem can be solved in many different ways, and just because you don't remember the exact way to solve it, you can still apply a variety of methods to arrive at the same answer, and 2) we can check our answer from the new method by checking it with the answer from an old method

i agree it's useful but i'm also tired of reading about drag and i did take 2.006

At least we end up learning about a tough topic in a much easier manner

Yeah I like the fact that we're tackling the same problems using different methods.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

I know this has been brought up before, but a short explanation of what the Reynolds number is (instead of your reader having to go to wikipedia) would go a long way.

We also went over this in class, which was very helpful. Again it might be nice to see this in writing

To prevent another student from having to go to wikipedia: its the ratio of inertial to viscous forces in a flowing fluid. It is used to distinguish between laminar (low Re, smoother flow) and turbulent flow (higher Re, characterized by more randomness, eddies, vortices, and other instabilities)

Thank you!

I think it would be helpful if you provided the equation of the Reynolds number here. I had forgotten the exact form of the equation we derived earlier by the time I read this memo.

since it's already been taught, i don't think it needs to be restated here. we could easily flip back and read about it again. save trees!

Oh I see, thank you for the post.

Why can't dim. analysis solve this?

Because it can never find those factors like "1/3" that we saw in the last section. it only deals with the groups of variables.

I think easy cases get us more accurate answers than dimensional analysis as it brings in the factors we had originally ignored (like the factor of 1/3 mentioned above).

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

I believe that, in class, we were actually able to use dimensional analysis to establish (at least the basic form of) the function, f, as the relationship between the volume of a pyramid and its height/base.

Is the, perhaps obvious, reason dimensional analysis doesn't work here because the drag coefficient and Reynolds number are, by definition, dimensionless?

If that is the case, then should we consider the methods to come a sort of "Plan B", in general, for when dimensional analysis fails?

Dimensional analysis never finds f for you. even for the pyramid, dimensional analysis only gave us V/hb^2=f(b/h), and then we had to use separate geometric reasoning to find f=1/3. Dimensional analysis always deals with dimensionless quantities.

Here, Drag coefficient and Reynolds numbers are particular combinations of other variables that yield a net dimensionless products. (Re = density * velocity * diameter / viscosity, for example).

Dimensionless analysis has not failed, it's just harder in this case to determine the function f. We used the physical experiment of dropping the cones to find two data points and we interpolated in a previous reading, but we can't extrapolate just from these points because of the (not particularly nice) behavior of f.

That was a really good explanation.

I think people are right about having a brief reminder of what Reynolds number is because I am slightly confused reading this paragraph

This was the only method used when I learned dimensional analysis before, but it's obviously somewhat limited in that you need to be able to measure a couple examples before you can find the trend.

Is there a particular reason why the data was limited when using cones? What would be a ideal situation for testing this?

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

I don't remember this result from the cone dropping...I only remember getting a good estimate of their velocity ratio

Oh, I get what you mean. Perhaps explain more clearly that air is a Reynolds number generally between 2k and 4k?

Maybe this can be found from the cone example, but we just didn't explicitly go through this part in class?

It'd be nice to see an annotation of the previous section that this was solved in for easy reference (i.e.: (1.2))

obviously cannot be proven using experiments, so we'll look at end cases for very small and very large cones?

Isn't this the idea of easy cases, we are looking at the two ends to see how it will respond in extreme situations.

I would expect that the surface area of the cones becomes more important for high Reynolds number.

I think it might be useful to just reiterate and explain what a low or high Re actually means

and also maybe an example of the extreme cases of the Reynolds number

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: Re $\gg$ 1 (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: Re $\ll$ 1 (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

Because they describe the limits?

I think so. Measuring when Re is in the range 2000-4000 doesn't seem to provide any intuition about the behavior of the function f in either extreme.

Well, if you do the experiments, you're adding data points and more data points will always give you a more accurate map of f. However, easy cases look at the extremes which should give a pretty decent map of f, but may be prone to error since it gives trends.

I think that he means that a map of f could be created by using a large range of Reynolds numbers, 'such experiments' referring not to only the extreme cases, but to experiments like the cone examples, but over a range of different Re's.

Also, how do you know what range is a limited range? How do you know that 2000 is not the minimum and 4000 is not the maximum limit. How do we know that the range goes from 0.1 to 10^6 and not from .00001 to 10^12?

You don't. You keep on doing experiments until you see little change, indicating that you're still in the same regime.

I would think that a high Reynolds number experiment on a paper cone would be nearly impossible

True, although you might be able to do it by dropping in an incredibly inviscid fluid instead of air.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{6.1}$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: Re ≫ 1 (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: Re ≪ 1 (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

Not especially, but I get what you're trying to say.

Depends on to which experiments are actually being referred. Conducting experiments where the Reynolds number is 10^6 would be rather difficult for the common student/reader.

Yeah, with a very high reyolds number, the paper cone experiment would be tough.

I dont understand why.

Because Reynolds number is pvL/u. So to achieve large Reynolds numbers, you will need extremely large density, velocity, or length (or even some large combination of the three!) parameters. The magnitudes of which are so large, most people would be unable to actually develop an experiment that would allow them to test it.

Probably at least more difficult than estimation...

Maybe just saying they are often not feasible given the average person's available materials - whereas anyone can take out a pen and paper and do the math.

That's a good thing to point out: the extra data points may give you more accuracy, but not necessary more accurate reasoning.

I thought we already knew the Reynolds number and it's implications. It seems to me that it would explain quite a lot.

It will be very helpful to have a diagram of a cone in this portion. It would definitely allow me to conceptualize everything better.

Wow, this reminds me of those recursive problems earlier in the book. (in particular the one with the "game")

Could you elaborate on why this reminds you of the recursive problem? I don't see the connection. Also, what section is the "game" problem you mentioned?

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

I know this was said earlier in another comment, but a general explanation of the Reynold number would be nice. Not only do I not know anything about the Reynolds number, but the following discussion about the meanings of different Reynolds numbers is lost on me.

A lot of what is needed to know about the Reynold's number was explained in earlier sections. It would be redundant to mention it again.

I think the issue is that the readings are spread out over the course of a semester, whereas if it was an actual paper book it would be easy to flip back to the section where it was defined.

Square or cube??

Good point, the readings about the pyramid first talks about the square base of the pyramid and then it goes into combining the 6 into a cube, whatever case is being referred to should be specified.

Does it really matter? We get the point–it's symmetric.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{6.1}$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

For the pyramids, it was quite intuitive as to what would constitute the "easy case". Here though, I can't say that it would occur to me to claim that Re >> 1 is an easy case. What is the reasoning process here?

He does say that "The physical reasoning in this regime is the subject of Section 6.3.1," so maybe you'll find your answer there.

I think "easy" refers to things like making things extreme or equal to some know quantity or solution (like in the atwood machine where we made one mass huge and then when we made the two masses equal). In this case, the "easy" part comes from taking"extreme" cases

I think people are reading a bit too far into the use of '1'. Re>>1 and Re<<1 simply mean that Re is relatively huge and relatively small, respectively.

What's important is that for our Re>>1 case, viscosity is negligible, and for Re<<1, viscosity dominates.

In general, "huge" for Reynolds numbers can be much larger than thousands or tens of thousands or more, depending on the geometry.

"Small" can mean anything from much less than 1, to "few", to less than a thousand, etc.

I guess someone who's taken thermo would find this fairly obvious, because the cases you're generally looking at are Re >> 1 and Re << 1 (and for Re 1 things get tricky).

is this doing the same this as the last reading with the pulley? examine two cases and extrapolate from there?

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

More explanation? I don't understand why there are two cases

just a quick thought..Re could be confused as being the Real part of the Reynolds number...even though it's not possible for it to be imaginary

The reynolds number is a ratio of inertia forces to viscous forces. Low reynolds numbers, especially less than 1 are dominated by viscous forces, meaning that the inertia forces (or density * velocity * characteristic length) are not important. I don't know much about these low Reynolds number scenarios, but I think these situations are for micro-organisms and probably many other interesting motions. And so having two cases is important since the motion is dominated by very different forces.

what would be the information lost in this case?

Eek, this is quite difficult to understand without first researching the Reynolds numbers. Granted, this was explained in lecture in class, but as a textbook, for students / instructors that may not have had the corresponding background lecture, this chapter would be confusing. It would be a good idea to have an explanation, especially since this will be compiled into a book.

I'm not sure where else this has come up, but I like that you actually reference the future sections in the intro. It gives the reader a heads-up

Is this saying there is an almost infinite range of the Re?

so under what condition do you assume Re<<1?

Maybe I just don't remember seeing this before, but it seems strange to reference these sections when they're about to be introduced to the reader for the first time.

I think that is a pretty standard text book technique "we will talk more about this in section ... "

In my experience, turbulent flow means a Reynolds number > about 2300. Why did you choose to simplify it to >>1 and <<1?

I think 1 is more for the low Re limit, in which Re <1. In the high Re limit, we're not thinking about laminar/turbulent transition yet, just about very very high Re and the resulting relation between drag and viscosity in that situation.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

People's comments about a definition of Reynolds number hold very true here, seeing at least the equation would make this statement more clear to the reader.

At what point in terms of Reynolds number does the flow officially become classified as turbulent?

And also, how do we know that it does?

Perhaps some intuitive explanation would help.

There's a period of transition Reynolds numbers between 2300 and 4000, but it's safe to say anything higher is classified as turbulent flow. In any case, it can just be assumed to be a high number over 10^3.5.

2300 is the accepted transition between laminar and turbulent flow. Very few flows are actually laminar, as it turns out.

Is there a background or reason for this cutoff?

DO we know this from experimental observation or from physical/theoretical law?

Is there a method to approximate this value as well or is there an actual value? In thermo, I learned that the transition period is between 2300 and 4000 which seems like a huge range?

I really like the description and pictures we had of this in class, it was nice to really see how drag worked visually

I agree. I also think it would be beneficial to describe what turbulent flow actually looks like and contrast that to uniform flow.

sad I missed that lecture :(

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \tag{6.1}$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

make the air thin like at high altitude

Can air be viscous? I thought this was more talking about fluids, so if v=0 then the object is in air

Air is a fluid. It has a viscosity 2 orders of magnitude less than water.

just to further clearify, fluids are things that deform under shear forces, so both liquids and gasses are fluids.

Or we could just all breathe in deeply at the same time...

this font similarity for v's is getting a little too intense

not sure if it was mentioned previously but mu is also used to represent viscosity, that might be a nice alternative.

since the Reynolds number is a function of many variables are you just choosing viscosity to be the independent variable

An equation of the Reynolds-number formula would be nice here.

Yeah, I understand that viscosity is inversely porporitonal to the reynolds number, but I have no idea where else the reynolds number comes from

Re = density * velocity * characteristic length / viscosity. It's a ratio of inertial forces to viscous forces.

It would be nice to see the equation so one can clearly see that viscosity disappears in the limit of high Reynolds number

Agreed. I have no idea where the Reynolds number comes from, and actually being able to see the equation and refer back to it in sections like this would be nice.

I would agree as well. For someone who hasn't internalized/memorized the Reynolds number equation, it would be nice to see the equation each time you bring it up.

I agree as well, it's a little more tiresome flipping back through the previous chapters.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

isn't viscosity the only term in the denominator of the Reynolds number? So if it went to 0, the Reynolds number would go to infinity. I'm not sure how it just "disappears".

It's not the Reynolds number that disappears, it's the viscosity that goes to 0 and thus disappears, or in other words, if there is drag, it doesn't depend on viscosity at that limit.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $Re \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $Re \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

is there drag force without viscosity?

only in a world where everything is ideal. viscosity is basically another word for friction.

There is absolutely drag force without viscosity, since the object still has to deflect the fluid in its way. The model Fd rho*v^2*A still applies.

We can apply this rule to objects in water too right? I believe they have viscosity as well.

But wouldn't it still approach zero as the the viscosity went to zero? I guess there is still some force... like to change the momentum of the fluid since it still has mass..?

I like the earlier comparison to friction...there is still drag on a frictionless surface, if I remember correctly. We tend to assume there is no friction due to air, but there is drag.

Remember that viscosity is a property of the fluid, not a property of the object.

Friction's 'purpose' is to create the laminar boundary layer around the object, 'enabling' the fluid to exert a viscous drag force. (When actually using Navier Stokes, etc., to model the flow you assume that the fluid in contact with the object is not moving relative to the object. Someone correct me if I'm wrong on this.)

But, regardless, there are inertial drag forces, as well, and that has to do with the object having to displace an amount of fluid per unit time proportional to v*A.

It is the relative importance of these two sources of drag force (inertial and viscous) that the Reynolds number measures. Re (inertial force)/(viscous force).

(To be a bit more technical:) Inertial force is F rho*v^2*D^2. Viscous force is F mu*v*D, where mu=rho*nu. Their ratio is (rho*v^2*D^2)/(mu*v*D) = rho*v*D/mu, which is the commonly presented expression for the reynolds number.

6.3 Drag

Although the method of easy cases is a lossy method – it throws away information – it gave us exact results in the preceding examples. Let's sharpen the easy-cases tool by applying it to that favorite difficult problem: drag. The conclusion of dimensional analysis was

$$\text{drag coefficient} = f(\text{Reynolds number}). \quad (6.1)$$

The remaining problem, which dimensional analysis could not solve, is to find the function f.

One approach to finding f is experimental. Drop cones of different sizes, use the geometry and terminal velocity to compute the drag coefficient and Reynolds number, and plot the results. We used this approach with two cones, finding that the drag coefficient was the same at a Reynolds number of 2000 and 4000. These two data points are only over a limited range of Reynolds number. What happens in other cases, for example when the Reynolds number is 0.1 or 10^6?

Such experiments would provide the most accurate map of f. However, these experiments are difficult, and they do not help us understand why f has the shape that it has. To that end, we use physical reasoning using the method of easy cases. When we applied easy cases to the pyramid, we chose h and b to make an easy pyramid (one that could be replicated and combined into a square). For drag, we choose the Reynolds number to simplify the physical reasoning. One choice is the regime of large Reynolds numbers: $\text{Re} \gg 1$ (the two falling cones are examples). The physical reasoning in this regime is the subject of Section 6.3.1. The other easy case is the regime of low Reynolds numbers: $\text{Re} \ll 1$ (Section 6.3.2).

6.3.1 Turbulent limit

When the Reynolds number is high – for example, at very high speeds – the flow becomes turbulent. The high-Reynolds-number limit can be reached many ways. One way is to shrink the viscosity ν to 0, because ν lives in the denominator of the Reynolds number. Therefore, in the limit of high Reynolds number, viscosity disappears from the problem and the drag force should not depend on viscosity. This reasoning contains several subtle untruths, yet its conclusion is mostly correct. (Clarifying the

what do you mean by "mostly" correct?

I think it's on the same order of magnitude but i am not sure

I would love to see an example of how easy cases are actually lossy, since this is the first type of approx. that is technically lossy.

I think that restating how the viscosity affects the Reynolds-number would be very helpful here.

This is a strange to put this sentence.

what are these untruths?

This is a sentence that I hope is explained later on or in class, because it's quite a tease.

what did you throw out at unimportant? and what do you mean by "mostly correct"?

subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 30].)

In other words, f is constant! The consequence is

$$F_d \sim \rho v^2 A, \quad (6.2)$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_d \equiv \frac{F_d}{\rho v^2 A} \quad (6.3)$$

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number).

Object	c_d
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [20], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small (Re $\ll$ 1):

$$c_d = f(\mathrm{Re}) \qquad (\text{for Re} \ll 1). \quad (6.4)$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \quad (6.5)$$

nice to see this mentioned. was worried that the generalizations made earlier in the paragraph would have been left without mentioning that there is much more to it than just high speeds leads to turbulent flow.

How did we arrive to f being constant?

yeah i'm a little confused here

What f and why is it constant? I agree that viscous forces are constant at zero, but I'm not sure how you can say that drag force is constant, or at least I don't see in what way it's constant.

Agreed, I don't understand how this conclusion was reached. I understand what was said in the previous paragraph, but I don't understand how you can make the jump to this.

I'm a bit confused too... You're saying it's constant with respect to the Reynold's number, but that is a function of velocity as well, isn't it? Or are we assuming velocity is constant also?

Well the previous paragraph talked about the drag force, so perhaps f is Fd?

I don't think it's directly saying that drag force is constant but rather that because viscosity drops out, f(Re) constant. Therefore using equation 6.1, the drag coefficient = f(Re) = constant.

I think that if you say that the function "f" is constant it will be a lot less confusing than just saying that f is constant. Saying f is constant makes it sound like some variable in an equation while, here, we're talking about the equation/function itself.

On second thought, I'm not really sure I know what is meant by f is constant. What does it mean for a function to be constant? I think understanding this could help me understand how we got to the conclusion that f is constant.

We got to the idea that the viscosity is 0 because we wanted a way to get an infinitely high Reynolds number. We then used this conclusion to eliminate viscosity from the problem. This only works if you continue to assume that the Reynolds number is infinite. So f(Re) essentially becomes f(infinity), which I suppose only makes sense if it's a constant function, not really dependent on the Reynolds number. But the whole thing is extremely poorly worded.

subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 30].)

In other words, f is constant! The consequence is

$$F_d \sim \rho v^2 A, \tag{6.2}$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_d \equiv \frac{F_d}{\rho v^2 A} \tag{6.3}$$

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number).

Object	c_d
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [20], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small (Re $\ll$ 1):

$$c_d = f(\text{Re}) \qquad (\text{for Re} \ll 1). \tag{6.4}$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \tag{6.5}$$

This seems like a pretty abrubt conclusion...

yeah I agree...is it possible to give some additional background on how this conclusion was reached?

i think it sounds abrupt because it's right after a very formal parenthetical. i think the transition could be smoothened

I agree. I have no idea where this came from.

This kind of seems like it came out of nowhere. Maybe you should mention the equation explicitly somewhere in one of these paragraphs so we can remember it from the last reading?

It's still really unclear to me how the drag force remains when the viscosity goes to zero...

what is the physical basis for the drag constant when there is no viscosity?

I think viscosity is basically how "thick" a fluid is. like how water is more viscous than syrup (it's more free-flowing). but even if something isn't thick, when an object travels through the fluid, it still comes into contact with the fluid, so there is still drag.

Why doesn't 0 describe a vacuum state? It seems to me that the lowest you could ever possibly get is no contact between an object an its medium of transport (i.e. a vacuum). Instead, it seems that 0 describes a superfluid, which seems to be something altogether different.

in 2.006, we denote crossectional area as Ac (A sub c) to remind us, might be helpful

How is cross section defined. I think the shape of a rocket has smaller Reynolds number than the shape of a car at the same velocity, even if they have the section area. Am I right? If this is true, why?

Well right below this it says the value depends on the shape of the object and how streamlined it is.

I think the cross sectional area in this example is the area perpendicular to the direction of motion.

this reasoning is super confusing

doesn't this just always = one? or am I mixing up v's?

subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 30].)

In other words, f is constant! The consequence is

$$F_d \sim \rho v^2 A, \quad (6.2)$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_d \equiv \frac{F_d}{\rho v^2 A} \quad (6.3)$$

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number).

Object	c_d
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [20], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small (Re $\ll$ 1):

$$c_d = f(\text{Re}) \qquad (\text{for Re} \ll 1). \quad (6.4)$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \quad (6.5)$$

I'm a little lost. I know we had previous readings on this drag coefficient but how is the c_d so small when you make the viscosity 0?

The c_d does not get smaller when viscosity gets smaller, it is solely dependent on the shape of the object. It quantifies how "aerodynamic" a shape is.

what makes you say it's "so small"?

How do we know what the force is without knowing the coefficient?

We are simply defining the drag coefficient. In order to find the value of the drag coefficient, we will have to measure the force (or vice versa, to find the force, we will have to know the drag coefficient).

I like this chart, but it may be helpful to mention at about what Reynolds number these are taken at. (What value constitutes a high Reynolds number)?

Yeah, I'd like some base values as well to judge these numbers from.

I like this table too. I'd like some more object examples to go with it, I find this very interesting.

And by more examples, I mean lower numbers to go with the table below.

yeah this table was super helpful to help me visualize how different shapes/objects are related to their drag coefficient and how the coefficients compare to each other.

Very helpful to see these in a chart, for quick access.

Wow, we do know this is true, but cool that we could reason it out so easily.

yeah, it was really neat to see a previous concept so easily come across here

Is there any cd that is greater than that of a flat plate? And if so how high is the highest cd?

I know we've showed this before but it still surprised me to see how it only depends on shape. The table really helped to drive this home

Well, after what we've gone through in class, what else COULD it depend on? Mass and weight don't make very much sense, neither does density.

subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 30].)

In other words, f is constant! The consequence is

$$F_d \sim \rho v^2 A, \tag{6.2}$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_d \equiv \frac{F_d}{\rho v^2 A} \tag{6.3}$$

Object	c_d
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number).

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [20], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small (Re ≪ 1):

$$c_d = f(\text{Re}) \qquad (\text{for Re} \ll 1). \tag{6.4}$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \tag{6.5}$$

Can we get a little more here? I'd like to see it done out, maybe as an exercise, but I was a bit lost in this section (at least on first read).

It sounds like these examples are more common everyday (ions in seawater, bacteria swimming- these things are basically happening all the time) but the high reynolds- number flows are definitely more common in experience... interesting

Yeah, it's bit curious if you think about relatively which is more common. Though it's more a matter of scale: both of these cases happen a LOT.

No, when you actually talk about fluid flows, turbulent flow is far more common. I'm not sure bacteria can be defined as a "fluid flow."

Would water droplets forming clouds fall into this category as well?

What are some other characteristics of Reynolds number?

What do you mean by this? We already know what Re describes.

Now I see the reason for the two cases is because different forces have more impact based on size, so it really is two different problems.

fog droplet falling? this sounds odd...

yeah, what's a droplet of "fog"?

never thought of that as related. cool

How is it related?

Could that be discussed briefly in class tomorrow?

Nevermind. I read the last paragraph.

Maybe add something here about how this will come up later?

are low Reynold's numbers also characterized by shape? fog droplets and bacterium kind of have the same shape...

I'm excited to read this section coming up!!

subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 30].)

In other words, f is constant! The consequence is

$$F_d \sim \rho v^2 A, \tag{6.2}$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_d \equiv \frac{F_d}{\rho v^2 A} \tag{6.3}$$

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number).

Object	c_d
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [20], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small ($\text{Re} \ll 1$):

$$c_d = f(\text{Re}) \qquad (\text{for Re} \ll 1). \tag{6.4}$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \tag{6.5}$$

Does the << just mean much smaller?

Yes, usually by several orders of magnitude.

Yes.

I've lost sight of the lesson we are learning in this section. What is the easy-cases tool we are sharpening?

I'm also a little confused on this. I think the easy-cases might be the extreme cases (where Re >> 1 or Re << 1) which we are trying to generalize to find f, but I'm not sure.

The easy cases we're looking at are the two limits. when Re is very large and when Re is very small. Thus, we are "throwing away information" –everything in between. But we know that in order for us to find a right answer, it must work for all cases. so we can mess around with these two cases to find a solution that can be generalized

I would have liked to see this a little farther up, just as a reminder.

subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 30].)

In other words, f is constant! The consequence is

$$F_d \sim \rho v^2 A, \tag{6.2}$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_d \equiv \frac{F_d}{\rho v^2 A} \tag{6.3}$$

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number).

Object	c_d
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [20], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small (Re ≪ 1):

$$c_d = f(\text{Re}) \qquad (\text{for Re} \ll 1). \tag{6.4}$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \tag{6.5}$$

but not all the things in the table are round (although all the examples given are). How does this work for non-circular objects?

Usually, if something isn't round (at least for pipe flow), you use an approximation for the diameter, called the "hydraulic diameter" (this is for standard Reynolds numbers based on diameter, radius is half that, of course).

The hydraulic diameter is equal to 4*(cross-sectional area)/(wetted perimeter). The wetted perimeter is the perimeter of the cross-section that is in contact with the fluid. You can easily show that for a circle, the hydraulic diameter is equal to the actual diameter.

I'm not entirely sure if this applies outside of the pipe-flow scenario, but this is the general idea.

To clarify, the hydraulic diameter still falls under the realm of Reynolds number based on diameter. Because the Reynolds number is a dimensionless number which only classifies a flow, you can substitute any length for diameter so long as you then define that its the "Reynolds number based on [that length]." For example, Reynolds numbers based on length (such as over a plate) go turbulent around 5*10ˆ5, as opposed to Reynolds numbers based on diameter go turbulent around 2300.

The short story is that it's all just a matter of defining your parameter so it's clear what the Reynolds number is saying.

Same issue as last time with it being difficult to distinguish v from nu. Heads up, everyone.

Could one of these variables be made bold or something to make them more distinguishable?

I agree. i had to double take

If the object is not spherical, how do we substitute this equation for this condition?

stupid grammar comment: for parallel structure, might just want to add a verb aka "reduce the object's speed"

Alternately, I see how "make" is also applied to "the object's speed low," so just the last part could be changed to "or the viscosity of the fluid high." for a parallel construction.

How high? Would ketchup work?

I like this clarification- it really makes sense

subtleties required two centuries of progress in mathematics, culminating in singular perturbations and the theory of boundary layers [6, 30].)

In other words, f is constant! The consequence is

$$F_d \sim \rho v^2 A, \quad (6.2)$$

where A is the cross-sectional area of the object.

Therefore, the drag coefficient

$$c_d \equiv \frac{F_d}{\rho v^2 A} \quad (6.3)$$

is a dimensionless constant. The value depends on the shape of the object – on how streamlined it is. The table lists c_d for various shapes (at high Reynolds number).

Object	c_d
Car	0.4
Sphere	0.5
Cylinder	1.0
Flat plate	2.0

6.3.2 Viscous limit

Low Reynolds-number flows, although not as frequent in everyday experience as high-Reynolds number flows, include a fog droplet falling in air, a bacterium swimming in water [20], or ions conducting electricity in seawater (Section 6.3.3). Our goal is to find the drag coefficient in such cases when Re is small (Re ≪ 1):

$$c_d = f(\text{Re}) \qquad (\text{for Re} \ll 1). \quad (6.4)$$

The Reynolds number (based on radius) is vr/ν, where v is the speed, r is the object's radius, and ν is the viscosity of the fluid. Therefore, to shrink Re, make the object small, the object's speed low, or use a fluid with high viscosity. The means does not matter, as long as Re is small, for the drag coefficient is determined not by any of the individual parameters r, v, or ν but rather only by their combination Re. So, we'll choose the means that leads to easy physical reasoning, namely making the viscosity huge. Imagine, for example, a tiny bead oozing through a jar of cold honey.

In this extremely viscous flow, the drag force comes directly from – surprise! – viscous forces. The viscous force themselves are proportional to the viscosity ν. In fact, the viscous force on an object is given by

$$F_{\text{viscous}} \sim \text{viscosity} \times \text{velocity gradient} \times \text{area}, \quad (6.5)$$

really effective example

Interesting visual.

I agree. Its a good example of an easy/extreme case. I think the format of this section (comparing extremely low and high Reynolds number) illustrates the point really well.

yep great example to help the reader visualize something with a huge viscosity!

It's a very vivid visual and something we've all had experience with.

an explanation would be helpful for me

Look at the example given in the paragraph before about a tiny bead oozing through a jar of cold honey, that should clarify what viscous forces are.

cute

interesting

Is there any way to quickly justify this, or do we just need to accept it for now?

You can always check the units.... if that works, it's at least mildly justified.

where velocity gradient is the rate of change of velocity with distance (so if the velocity does not vary, then there is no viscous force), and the area is the surface area of the object. Because the drag is due directly to viscous forces, the drag force is also proportional to viscosity:

$$F_d \propto \nu.$$

This constraint is sufficient to determine the form of the function f and therefore to determine the drag force. Start with the result from dimensional analysis:

$$\frac{F_d}{\rho_{fl} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The viscosity ν appears only in the Reynolds number, where it appears in the denominator. To make F_d proportional to ν requires making the drag coefficient proportional to Re^{-1}. Equivalently, the function f, when $Re \ll 1$, is given by $f(x) \sim 1/x$. For the drag force itself, the consequence is

$$F_d \sim \rho_{fl} r^2 v^2 \frac{\nu}{vr} = \rho_{fl} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanic must do a messy and difficult calculation that is possible only for a few special shapes. For a sphere, the British mathematician Stokes showed that the missing constant is 6π; in other words,

$$F_d = 6\pi \rho_{fl} \nu v r.$$

This result is called Stokes drag. In the next section, we will use this result to study electrical properties of seawater.

Isn't this acceleration?

This is pretty intuitive, isn't it?

It is, but it may not be obvious when reading. I think this is very important stated here. Thanks for it!

how is it sufficient? and why didn't we know this 2 pages ago?

It seems a lot of "easy cases" are linked to dimensional analysis in some way.

You probably have defined the variables before, but please define them again. A disadvantage of an e-book is that it is difficult to look things up.

I really like the tables that show variables' meaning, just like in the last sections.

Perhaps a table at the beginning (or end) of every chapter with all of the variable definitions used in that chapter would be helpful – that way we would only need to refer to one particular page, rather than looking through the entire text. Things like the formula for Re could also possibly be included (but not derived there), just so that it's easier to find for later reference.

what was fl? flat plate?

Or maybe a page of formulas that is added to each we learn something new.

Although I agree with this as he discussed in class I'm not sure how this could be easily fixed, maybe make one bold? I don't think it's too big of a problem if you've been following the previous steps though

Once again, the nu and the v look very similar, and it's still quite confusing.

What if we used mu=nu*rho in place of nu, throughout? (dynamic instead of kinematic viscosity?) The math works the same, the reynolds number is more familiar (at least to course 2-ers, i think), and we avoid the tricky v-nu similarities.

I think this is a great idea.

Quick review of f(vr/viscosity). The f() does not change the units at all right? So then we get F_d/pr^2v^2 through dimensional analysis? Also, are we just supposed to know that f(x)=1/x when the Re<<1?

where velocity gradient is the rate of change of velocity with distance (so if the velocity does not vary, then there is no viscous force), and the area is the surface area of the object. Because the drag is due directly to viscous forces, the drag force is also proportional to viscosity:

$$F_d \propto \nu.$$

This constraint is sufficient to determine the form of the function f and therefore to determine the drag force. Start with the result from dimensional analysis:

$$\frac{F_d}{\rho_{fl} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The viscosity ν appears only in the Reynolds number, where it appears in the denominator. To make F_d proportional to ν requires making the drag coefficient proportional to Re^{-1}. Equivalently, the function f, when $Re \ll 1$, is given by $f(x) \sim 1/x$. For the drag force itself, the consequence is

$$F_d \sim \rho_{fl} r^2 v^2 \frac{\nu}{vr} = \rho_{fl} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanic must do a messy and difficult calculation that is possible only for a few special shapes. For a sphere, the British mathematician Stokes showed that the missing constant is 6π; in other words,

$$F_d = 6\pi\rho_{fl} \nu v r.$$

This result is called Stokes drag. In the next section, we will use this result to study electrical properties of seawater.

I had to re-read this a few times understand what was being said. It was not apparent what was going on

This might have been covered earlier, but why do we want to make Fd proportional to v?

Two paragraphs earlier, we found that for low Re flows, drag force is proportional to viscosity.

Thanks, this helped me too!

This might be extremely nitpicky, but here you have two different ways of expressing 1/something. Shouldn't you just pick a convention and stick with it?

Why is the function f given by this for low Re? Could someone please elaborate on this, I'm a little confused.

why didn't we look at Re>>1?

Yeah I was waiting for that case to be explain after R <<1

We did look at Re >>1; its the section before - the turbulent case.

Er, we did. Re>>1 is the turbulent case - the first one discussed.

I'm a little confused on how these constants are in the equation.

In the equation prior to the above paragraph, it gives Fd/(this expression) = f(vr/v). Solving for Fd gives this expression in front.

I think you've lost a factor of r here.

It's also interesting that nothing else about the material matters...it really is the fluid at low Re

This equation does seem dense.

so did this get you any farther than where you started? I feel like were in the same place- not completly knowing the relationship all we did was move some variables around

where velocity gradient is the rate of change of velocity with distance (so if the velocity does not vary, then there is no viscous force), and the area is the surface area of the object. Because the drag is due directly to viscous forces, the drag force is also proportional to viscosity:

$$F_d \propto \nu.$$

This constraint is sufficient to determine the form of the function f and therefore to determine the drag force. Start with the result from dimensional analysis:

$$\frac{F_d}{\rho_{fl} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The viscosity ν appears only in the Reynolds number, where it appears in the denominator. To make F_d proportional to ν requires making the drag coefficient proportional to Re^{-1}. Equivalently, the function f, when $Re \ll 1$, is given by $f(x) \sim 1/x$. For the drag force itself, the consequence is

$$F_d \sim \rho_{fl} r^2 v^2 \frac{\nu}{vr} = \rho_{fl} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanic must do a messy and difficult calculation that is possible only for a few special shapes. For a sphere, the British mathematician Stokes showed that the missing constant is 6π; in other words,

$$F_d = 6\pi \rho_{fl} \nu v r.$$

This result is called Stokes drag. In the next section, we will use this result to study electrical properties of seawater.

Maybe I missed something but when did this missing constant come into play? Also it would have been great for a last paragraph summarizing our methods of using easy cases.

We were originally solving for the dimensionless drag coefficient c_d, but it looks like we found F_d instead..I guess I'm kind of confused as well.

Same here, I think a summarizing paragraph to wrap up this section is really needed. It says dimensional analysis alone is insufficient which I agree with and understand but wasn't the point of easy cases to help us finish what dimensional analysis couldn't, or are you showing here that there are cases like with drag where you need complicated equations to actually figure out the missing constants.

Yup, that would be useful. This isi a realtively short chapter so it's won't be ridiculous to add to the end of it.

yeah, that sentence directly after the above epxplanations and equations kindof confused me about the poin of easy cases.

I totally agree. I was really confused after reading this section, even after rereaing for the 3rd time! What I think is going on is that we said (from dimensional analysis) that Fd/rho_fl*r^2*V^2) = function of Re. The missing coefficient should represents that "function" of Re. And even after doing dimensional analysis, we still can't find it. So, my conclusion is that he's concluding that this is a case when an "easy case" isn't so easy.

I'm confused by this. I thought we got the answer using dimensional analysis, but this is now telling us that we didn't?

Yeah this is also a little unclear to me.

The easy case turned into no longer easy at this point for me.

where velocity gradient is the rate of change of velocity with distance (so if the velocity does not vary, then there is no viscous force), and the area is the surface area of the object. Because the drag is due directly to viscous forces, the drag force is also proportional to viscosity:

$$F_d \propto \nu.$$

This constraint is sufficient to determine the form of the function f and therefore to determine the drag force. Start with the result from dimensional analysis:

$$\frac{F_d}{\rho_{fl} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The viscosity ν appears only in the Reynolds number, where it appears in the denominator. To make F_d proportional to ν requires making the drag coefficient proportional to Re^{-1}. Equivalently, the function f, when $Re \ll 1$, is given by $f(x) \sim 1/x$. For the drag force itself, the consequence is

$$F_d \sim \rho_{fl} r^2 v^2 \frac{\nu}{vr} = \rho_{fl} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanic must do a messy and difficult calculation that is possible only for a few special shapes. For a sphere, the British mathematician Stokes showed that the missing constant is 6π; in other words,

$$F_d = 6\pi\rho_{fl} \nu v r.$$

This result is called Stokes drag. In the next section, we will use this result to study electrical properties of seawater.

What are the other shapes, out of curiosity?

I'd guess the ones seen in the table above but I'm not sure

Except for the car...

yup

or the ones with defined/simple boundary conditions (haha basically the ones that we would do in a fluid mechanic class, but in the real world, i guess a lot of conditions are much harder to define so you guesstimate a lot)

Did he reason this mathematically or experimentally (or both)?

It seems like we jumped to this, our methods got us most of the way but I'm curious as to how exactly he came up with this number

I would also be interested in knowing how he came up with this value, as a comparison with the approximation methods we've used here.

Given he is a mathematician, I would say that he solved it by messy integrals of the unpleasant kind...

I think we need more here too. I think this section needs to be expanded.

referring to the Navier-Stokes Eqt?

so "stoked"

is this releated to the stokes shift?

That's cool, I like this example.

So do I, it's cool to learn the progression of how different scientific methods and discoveries came about.

Would we approximate the 6 pi?

I think dimensionless consants in general are almost impossible to approxmate. For a sphere, you could've seen the pi coming, but that 6 is not as simple.

Agreed - the dimensions are a easier to come by.

I'm pumped to apply this to Navier-Stokes.

where velocity gradient is the rate of change of velocity with distance (so if the velocity does not vary, then there is no viscous force), and the area is the surface area of the object. Because the drag is due directly to viscous forces, the drag force is also proportional to viscosity:

$$F_d \propto \nu.$$

This constraint is sufficient to determine the form of the function f and therefore to determine the drag force. Start with the result from dimensional analysis:

$$\frac{F_d}{\rho_{fl} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The viscosity ν appears only in the Reynolds number, where it appears in the denominator. To make F_d proportional to ν requires making the drag coefficient proportional to Re^{-1}. Equivalently, the function f, when $Re \ll 1$, is given by $f(x) \sim 1/x$. For the drag force itself, the consequence is

$$F_d \sim \rho_{fl} r^2 v^2 \frac{\nu}{vr} = \rho_{fl} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanic must do a messy and difficult calculation that is possible only for a few special shapes. For a sphere, the British mathematician Stokes showed that the missing constant is 6π; in other words,

$$F_d = 6\pi\rho_{fl} \nu v r.$$

This result is called Stokes drag. In the next section, we will use this result to study electrical properties of seawater.

I'm interested to see how this is used for this application.

Conditions where Re < 1 also lead to neat things like reversible flow, which is pretty uncommon in fluid mechanics studies. You could put a drop of dye in some viscous liquid, stir, then reverse your path and the dye would regain the original shape. I suppose this also depends on the diffusion time constant being relatively long, which is captured by another dimensionless parameter.

Again, I'm confused about what exactly is meant by "easy cases". I don't really know what methods were used to solve the problem here, and I certainly couldn't try to use this method to solve a problem myself.

The only easy case I caught was given right at the end for Re<<1, but I'm don't really even understand that one.

I agree, I don't think I understand the basic concept of easy cases, or how it is applied in this case. It just seems to be more dimensional analysis

I think a short "wrap up" paragraph at the end could go a long way in describing how the method was used, etc., and just overall clarification.

At the beginning of the section, I think Sanjoy defined an "easy case" to be a problem in which we have to sacrifice accuracy because the problem is so complex. However, in some of the first "easy cases," no accuracy was sacrificed. In this one, his definition holds. I guess I'm confused about this also, just in a different way.

Easy cases tend to be extremes. Like Re<<1 and Re>>1. Sometimes numbers like 0 and infinity are easy cases.

I was expecting you to show a comparison between these 2 functions and the actual behavior observed (the graph displayed in the last section) so we could see just how close our estimations got to the actual thing.

Yeah, I agree, I'd like to see a graph... I find I tend to lose some conceptual meaning to things when I'm just staring at equations.

That plot comes next (as we saw in class).

where velocity gradient is the rate of change of velocity with distance (so if the velocity does not vary, then there is no viscous force), and the area is the surface area of the object. Because the drag is due directly to viscous forces, the drag force is also proportional to viscosity:

$$F_d \propto \nu.$$

This constraint is sufficient to determine the form of the function f and therefore to determine the drag force. Start with the result from dimensional analysis:

$$\frac{F_d}{\rho_{fl} r^2 v^2} = f\left(\frac{vr}{\nu}\right).$$

The viscosity ν appears only in the Reynolds number, where it appears in the denominator. To make F_d proportional to ν requires making the drag coefficient proportional to Re^{-1}. Equivalently, the function f, when $Re \ll 1$, is given by $f(x) \sim 1/x$. For the drag force itself, the consequence is

$$F_d \sim \rho_{fl} r^2 v^2 \frac{\nu}{vr} = \rho_{fl} \nu v.$$

Dimensional analysis alone is insufficient to compute the missing magic dimensionless constant. A fluid mechanic must do a messy and difficult calculation that is possible only for a few special shapes. For a sphere, the British mathematician Stokes showed that the missing constant is 6π; in other words,

$$F_d = 6\pi \rho_{fl} \nu v r.$$

This result is called Stokes drag. In the next section, we will use this result to study electrical properties of seawater.

From this example, I don't know if this section should be called easy cases. Aren't they just complex cases that can be solved experimentally?

I think it would be a lot easier to do empirical studies of the mid-range of Reynolds numbers. Was it possible to solve the low or high Re case experimentally and easily? We've taken the problem of drag, which is complex overall, and limited our study to just the easy cases within it.

Overall, I thought this section was rather ironic: While it is based on easy cases, I found these examples to be some of the more complicated ones we've seen thus far in the course.

But these easy cases are still much easier than solving for drag along the entire range of Re using Navier-Stokes or constructing expermiental apparatus.

6.3.3 Conductivity of seawater

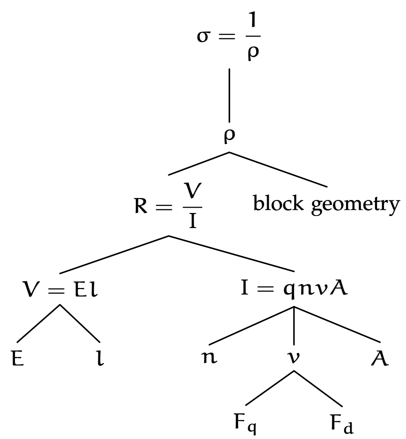

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

I thought this method of approximating the conductivity of seawater was difficult. 8.01 and 8.02 were sufficient to allow me to follow the math, but I'm not sure the level of detail used reflects what I would be able to think up on individual problems. I wish there were more frequent ties to the intuitions behind each assumption, as well as connections back to the original goal (how it fits in to the larger context as an example of an easy case).

As someone already said, it can be easy to get caught up in the detailed math and forget purpose/lose interest.

6.3.3 Conductivity of seawater

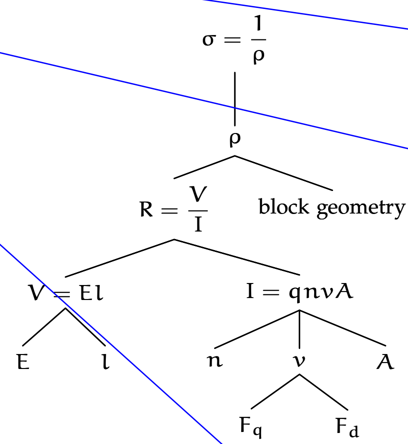

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

this saltwater problem seems more like divide and conquer than easy cases.

Read these two subsections for Tuesday's memo.

I thought this was a good example for us to walk through since it's so involved, but I'm very confused: how is this an easy case? In my mind, there are way too many calculations involved here for it to be an easy case, and we don't do too much simplification of the problem.

I agree. I'm not sure where we're headed with this example or what we're trying to show.

It felt that way to me in the previous section also. I don't think I'm fully grasping the concept of "easy case".

I think this is definitely a confusing example, but by easy case we're just looking at an extremely low reynolds number, one close to zero. Because these cases are hard to find in nature it involves a lot of calculations but in the big picture it is an easy case.

I don't think easy case means easy to see, but that with the right simplifications, it makes the calculations and evaluations easy.

I agree, this example is another great way to showcase the relevance of Reynolds numbers and real world cases where they can be very low.

Is this saying that situations with low reynolds numbers are rare or that conductivity in seawater is rare? It seems to me that conductivity in seawater is constitutive and so I would guess that the meaning is the former but could this be clarified in the reading?

I read it as referring to the rarity of low Reynolds numbers.

Particularly after the discussion in class regarding the difficulties in achieving low Reynolds numbers experimentally.

I also think this means that examples with low Reynolds numbers are rare.

Yeah, I agree with this - it makes more sense than the other interpretations of the sentence.

I like that things we've learned earlier keep coming back at a reasonable frequency. When I am about to forget, it gets integrated back to the new material.

6.3.3 Conductivity of seawater

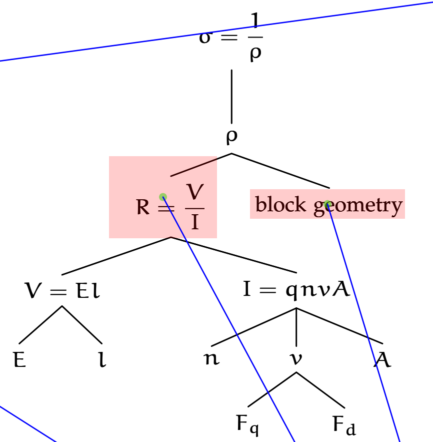

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

I'm a little confused, what exactly are we after in this example? The drag force imparted on a charge?

I didnt think about that until you mentioned it, but its a good point...charge flow? or something else? I guess we'll see!

to me, I was just thinking a more traditional approach to conduction (remembering that putting the hairdryer cord in the shower is bad...)

I'm a little confused also, with just a casual glance down the pages(because I thought reynolds number had to do with drag, I didn't see anything about drag or reynolds number). Maybe it would be helpful to very briefly state what the reynolds number will be used for?

The sense I got in lecture is that, given that current in seawater is primarily carried by the flow of Na+ and Cl- ions, rather than by lone electrons, finding the drag force on a single ion moving through the water will allow us to estimate the speed at which those ions flow given a particular electric field, which will tell us the conductivity.

I'm happy to see the return of the tree. We haven't used them in a while.

yeah this is a good place to refresh our memories about trees and the divide & conquer method.

This is a good refresher back to divide and conquer, and also shows us that many problems involve combination of methods we've learned so far.

we are back to unit 1

I always appreciate a good refresher...keep our D&C skills sharp!

Could it be that we're learning to solve these problems through an application of many different methods we've learned?!

is this the resistance?

Is this referring to something like the length and area of a block of seawater?

I think so

Maybe it's a control volume for later calculations.

yes it's mentioned later

6.3.3 Conductivity of seawater

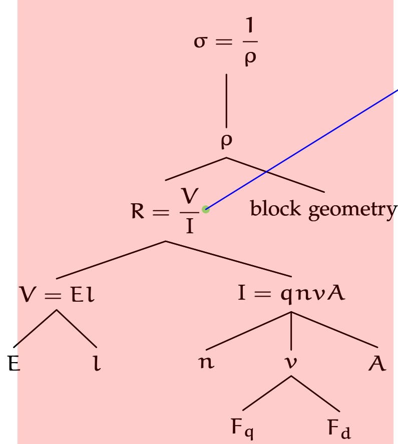

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

can you define these variables for us?

The variables are defined as you go keep on reading. Perhaps this diagram could be placed later in the section so that it actually makes sense?

I agree, I figured out what "solution tree" was being referenced, but glancing at it right now just made me more confused

This diagram seems a little too early. I didn't understand what it was for until the next page.

Also, it would be helpful if we numbered the different parts of the tree and had them correspond to the text where they are mentioned. That way we could more easily understand the problem.

I'm not sure thats a great idea...I think it would make things overly complicated. Placing the image lower in the text to allow the reader to catch up is definitely helpful though.

It would be helpful to have this tree after all the variables are defined. It would be easier to follow.

I actually think its fine where it is. you may not understand it at first but it is very clear as you go through the reading so in the end you understand it.

Just a thought, but perhaps you could put the questions in bold? Often times I'll start reading and get lost in the details so by then end I'll forget what we were trying to find in the first place. Like a diagram, it would help me stay focused on the main idea.

Interesting- I didn't realize it was the ACTUAL reciprocal

I think that it might be more helpful to define conductivity on its own terms, as opposed to define it relative to another term.

Well, it is logical since resistivity is a previously discussed concept.

You start to explain conductivity but more end up describing resistivity...maybe just put in a sentence to clearly define conductivity.

Conductivity is the inverse of resistivity, so he needs to describe resistivity to describe conductivity.

6.3.3 Conductivity of seawater

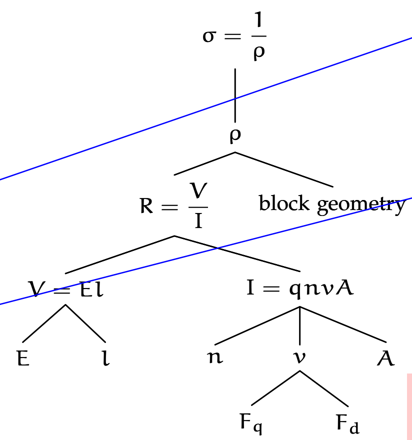

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

I'm not familiar with resistivity, could the term be quickly defined here or even below when resistance is defined?

rho = resistance * (Cross sectional area / length)

Instead of writing an apology here, can you just call resistivity another variable to not confuse it with density? rho sub r? phi? I think saying, we're going to call resistivity phi, as to not confuse it with other variables, and it has units of ohm*m would be good.

Having a subscript would be okay, but I don't think changing variables is a good idea. There are conventions, and then if you see something similar later, it is easier to relate if you stick to the conventions.

I agree with this. I'm already very confused with all the variables and equations in this course.

I feel like resistivity is normally rho and you don't have to apologize or change the variable. It's just convention.

I like how this question is addressed right after conductivity is mentioned, because it's something I asked myself as soon as I read it.

Perhaps this is a vote for early bird gets the worm, but it's really overwhelming to read these paragraphs with myriad grey lines overlapping.

It would be helpful to know how it is related.

I like that the things we're learning are tying together. The initial tree diagram on the first day of classes explaining what we were going to learn made them seem quite independent.

Isn't rho per unit length while R is a total amount of resistance?

I was under this same impression from 8.02 and the like.

I think it will be helpful to show a equation of their relation. It really isnt clear as written here.

6.3.3 Conductivity of seawater

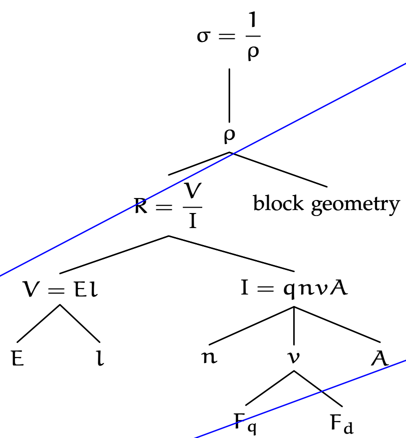

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

maybe reword as: "then what is the difference between rho and R?"

well i think the question is more about the advantages of having two different related conventions. so it's not so much how they're different. also, i'm glad he answered this question as i was curious

for that matter...why have rho and sigma? if one is just the reciprocal of the other, don't they tell us essentially the same thing...?

So, already knowing the difference between resistance and resistivity, this explanation actually left me more confused. It could be more straightforward... (i.e., maybe getting rid of "Resistivity is related... Then why have..." and just explaining one is based on the specific geometry in addition to the material properties.)

I think it could be changed slightly to say "... for a given wire material, the resistance depends on...". ("material" being the key addition).

What is an intensive quantity? and vs an extensive quantity?

An intensive property is one that doesn't depend on the physical size of the system or the amount of material in the system; an extensive property does depend on those parameters.

In that case, why does the text say that resistance depends on the wire's length and cross-sectional area, yet is not an intensive quantity? Isn't that contradicting?

..no, that's a direct violation of the definition of intensive. Resistance depends on length, area, and resistivity. Since length and area are shape-dependent, so is Resistance. So it is NOT intensive.

Resistivity, on the other hand, is a material property.

so I guess for a fixed size and shape the resistance is know and so it is "intensive" but if the total resistance is yet to be determined (depends on length and width) it is "not intensive"? Seems like with enough information you could turn something that has resistance that is "not intensive" into an "intensive" quantity.

Thanks for the definition!

Thanks for this - it really cleared up this paragraph. I think it would be useful to include this, maybe as a sidebar?

6.3.3 Conductivity of seawater

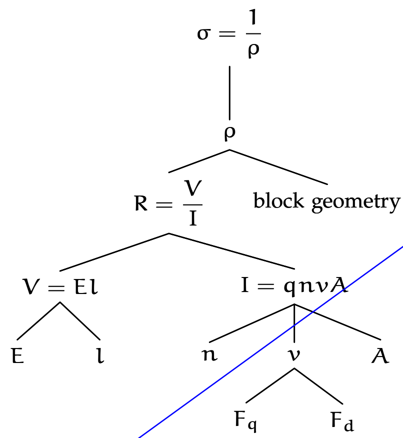

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

I don't see how these few sentences are relevant since we don't even know what intensive or extensive means yet. You should just sketch the solution tree and not start that little tangent maybe.

that teaser isn't fair! you have us curious, but aren't going to answer why it's not an extensive quantity.

What does it mean for a quantity to be intensive/extensive?

intensive variables don't depend upon size or weight, and therefore can be scaled. extensive variables depend on these properties and are case-specific.

I don't really like this sentence, it seems out of character for him to put something in technical jargon without explaining it.

Agreed, it would probably be better to put a simple definition intensive quantity in parens instead.

These are good terms to bring up though–very good to keep in mind when analyzing physical problems. They relate closely to problems of scale and dimensions that we've been working on.

Is it not extensive since even though it depends on size via cross-sectional area and length, it's not a direct relation with volume?

This is a really interesting comment. I think an appendix should be put in at the end of the chapter (a lot of books have answers to weird questions at the end of the chapter) where it shows why this is true. But I didn't know that; very interesting point. Makes sense.

6.3.3 Conductivity of seawater

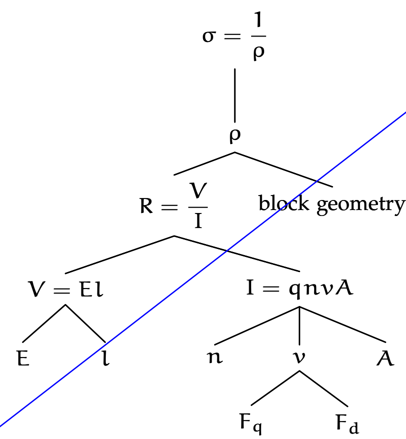

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

How can something be neither extensive nor intensive? Please explain.

I think its because it changes but isn't always reliant on the shape or size? I don't think that's true though, so I'm not sure

I agree. This is kind of a perplexing statement and I'd like to see some further explanation on how something can be neither intensive nor extensive. Maybe it just means it's sort of something in between?

what does intensive or extensive mean anyways?

I think it is essentially like the difference between density and length: one is constant for across any quantity of a given material and one is dependent upon physical alteration. This is just a guess, though.

This could definitely be made more clear. Perhaps even including the equation R = rho*d/A and explaining that rho is a property of a given material while R is a property of a given material and geometry would help clarify the difference

This little 2-sentence block of extensive vs. intensive would definitely benefit from some clarification...

does every material have a resistance and a resistivity? or are they mutually exclusive is resistance more of a material property and resistivity more dependent on the size and shape?

What about resistivity? Is it intensive, extensive, or other?

Maybe you should say "R times geometry" since nothing is actually added.

or just "R and geometry"

"and" is the clearest thing I can think of too

yeah i agree, the "plus" is confusing but an "and" would make the sentence much clearer.

Thanks for clearing this up. It was confusing.

6.3.3 Conductivity of seawater

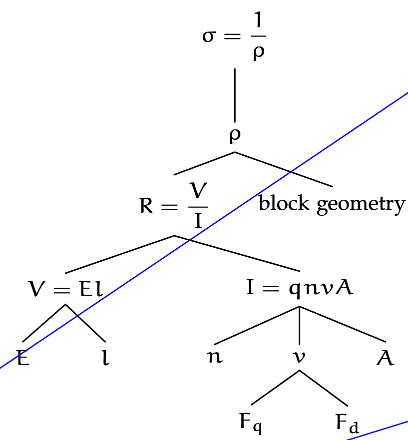

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

solution for what? i'm still confused as to where all this is going.

Agreed, I'm not sure what the point of all this is. I know you mentioned finding the conductivity of seawater, but why do we want to know this? I think a little explanation on why this calculation or result is useful would be helpful.

I agree too, little confused because it seems like conductivity is the purpose but resistivity formulas dominate the intro.

I've seen it used to find the salinity of the water (since that will change from location to location and conductivity is easier to measure). However, we use salinity to find conductivity in this case so I'm not sure why its important

This might be a better place to have the diagram or even just the top section of the diagram.

I agree, the diagram works well but not in the place it's at.

What exactly is a solution tree? I feel like this is a little shady, since afaik, this isn't a standard diagram or method, and it wasn't even mentioned where this solution tree was started...

It's weird to think of rho depending on R, instead of R depending on rho... Maybe a comment on why we are attacking the problem this way?

I think the reason we are doing this is not because it physically relies on R but instead because in our case, finding rho depends on already having R.

I think the coolest part of this class is that it integrates various disciplines... Physics, math, chemistry, and etc.

Are we just supposed to know this?

Well, knowing that resistivity is dependent on resistance is a matter of knowing the formula, and the extrapolation to voltage and current is pretty standard when it comes to the study of electricity (witness the entirety of 8.02). The next formula is pretty standard as well.

I like how this model is simple enough that anyone could come up with it (basic electricity knowledge from 8.02)

6.3.3 Conductivity of seawater

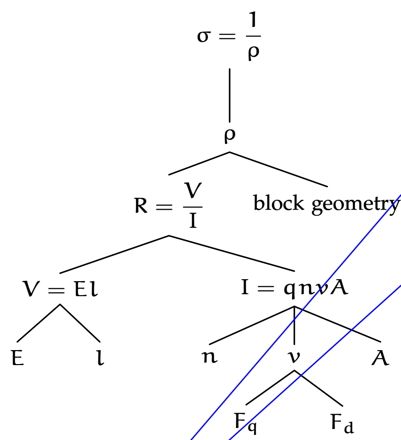

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

How do you isolate a block of seawater? Isn't there some other unit of seawater that is more appropriate? (that is applicable to liquids?)

I have a feeling we'll figure this out for the "block" and then extend it to a larger value?

So this explains why water and seawater have high conductivities?

Yes. In fact, pure water does not conduct electricity at all as there are no free ions to transport the charge.

Which, from chemistry, implies that the water we drink is not ion-free.

Kind of a random question..but does that mean if you were to use an electronic device in pure water you wouldn't get electrocuted?

I think so!

This read a bit awkwardly until I realized "transport" was the verb. Maybe say "it is the ions" or "...are what transport the charge"

I agree that this is confusingly worded. Perhaps the simplest helpful change would be to change "dissolving" to "dissolved"

Did you get this by dimensional analysis or is this a physics equation I've long forgotten?

As others ask, it would be good to say what the block is.

but if we are talking about the entire ocean how is there a cross sectional block?

but if we are talking about the entire ocean how is there a cross sectional block?

but if we are talking about the entire ocean how is there a cross sectional block? are we doing the calculation per unit volume?

So we are just assuming that all of the ions here are either Na or Cl? so Na+ and Cl- each have a charge of +1 or -1. and theoretically they would be present in equal amounts. wouldn't the + and - charges cancel each other out? or does conductivity only depend on the absolute value of the charge (the magnitude of the charge)

6.3.3 Conductivity of seawater

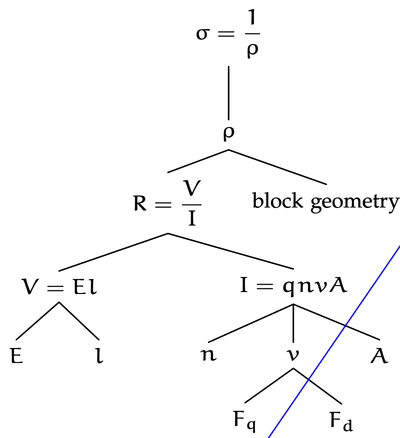

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

but if we are talking about the entire ocean how is there a cross sectional block? are we doing the calculation per unit volume?

Yes, otherwise the calculations wouldn't make any sense.

if you think about it, we are worried with the resistivity, which will take into account a certain volume and a certain number of things in that volume, so it more like density in that it doesn't matter that we are only using a block of ocean. Assuming that the ocean is uniform in salt concentration, then the whole ocean will have the same resistivity as any size block (its an intensive property).

It might be good to describe the block you're referring to. I guess it's a unit cube?

I think you can think about it this way: Resistivity is an intrinsic property of seawater, so we should always get the same value no matter what configuration of seawater we measure. The easiest way to get an accurate measurement here might be to place two plates of area A on either end of a uniform chamber filled with seawater. That way we know exactly what the relevant cross-section is and can measure the resistivity.

what do you take as the cross sectional block? are we doing the calculation per unit volume?

How do we determine this in a large body of water? Any current would take the path of less resistance right? and not use the full area?

Using the full area does lead to the least resistance. Charge in a wire doesn't just flow down an arbitrary middle core of metal ignoring the outer part of the conductive wire.

remember that is why we are considering resistivity and not resistance, I'm pretty sure we are assuming things to be uniform also

I am a bit confused by the description of these variables.

6.3.3 Conductivity of seawater

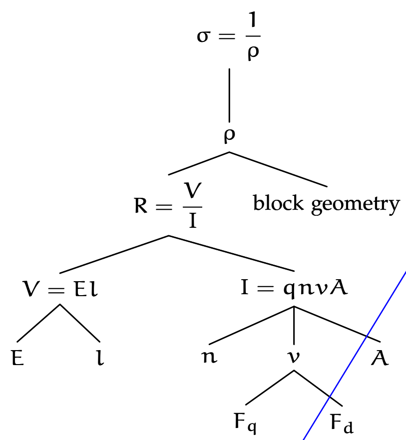

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

Not sure what it is, but I find this wording very strange. number of ions/m^3 or something along those lines would make it more clear.

I think he prefers to use dimensions over units, though. So perhaps ions/L^3, although I didn't find this confusing.

How do we calculate this?

You can look up the composition of seawater to find out the concentration of each ion.

This seems analogous to mass fluid flow, if ion concentration were density (dimensions of mass/volume) and the left side had dimensions of mass/time.

The "its" confused me because I didn't know what to which quantity "its" was referring.

Yeah I was about to say the same thing.

I think it's ions? Though, I'm not really sure either...

Yes, it is the ions, since moving charges create currents. However, the "its" is quite ambiguous.

I don't remember having to deal with terminal speed when calculating current, maybe a little elaboration?

maybe I'm misinterpreting this, but it seems like this should be the "concentration of the water (ions per volume)" because as it reads now, it doesn't make sense to me to think about concentration of the ion as "ions per volume"

I think it means that we are simply interested in the amount of ions in some arbitrary volume.

We aren't really interested in the concentration of the water, as water isn't the conductor.

The text is probably considering the quantity of ionic units in any volume so it's not really relevant what the concentration of water itself is.

can electricty travel slower than the speed of light?

Its good that we are using all these previous methods, but when are we going to see the easy case thing? I feel like we are just summarizing the semester with this problem so far.

6.3.3 Conductivity of seawater

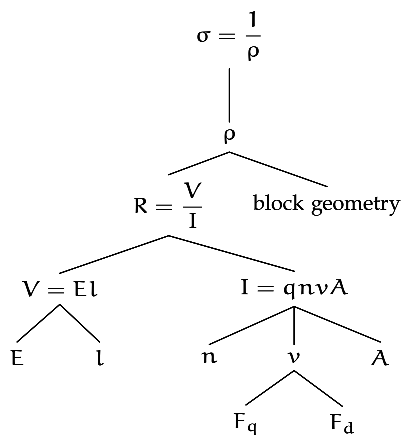

To illustrate a rare example of a situation with low Reynolds numbers, we estimate the electrical conductivity of seawater. Doing this estimate requires dividing and conquering.

The first question is: What is conductivity? Conductivity σ is the reciprocal of resistivity ρ. (Apologies for the symbolic convention that overloads the density symbol with yet another meaning.) Resistivity is related to resistance R. Then why have both ρ and R? Resistance is a useful measure for a particular wire or resistor with a fixed size and shape. However, for a general wire, the resistance depends on the wire's length and cross-sectional area. In other words, resistance is not an intensive quantity. (It's also not an extensive quantity, but that's a separate problem.) Before determining the relationship between resistivity and resistance, let's finish sketching the solution tree, for now leaving ρ as depending on R plus geometry.

The second question is: What is a physical model for the resistance (and how to measure it)? We can find R by placing a voltage V – and therefore an electric field – across a block of seawater and measuring the current I. The resistance is given by $R = V/I$. But how does seawater conduct electricity? Conduction requires the transport of charge. Seawater is mostly water and table salt (NaCl). The ions that arise from dissolving salt transport charge. The resulting current is

$$I = qnvA,$$

where A is the cross-sectional area of the block, q is the ion charge, n is concentration of the ion (ions per volume), and v is its terminal speed.

To understand and therefore rederive this formula, first check its dimensions. The left side, current, is charge per time. Is the right side also charge per time? Do it piece by piece: q is charge, and nvA has dimensions of T^{-1}, so $qnvA$ has dimensions of charge per time or current.

Good step to remember!

This is always a great check to make sure you haven't gone astray.

This seems like it was stated a bit hastily. I find it easier to understand the points about dimensions if they are in a table, or at least some consistent format throughout the entire book.

Does this have dimensions of T^-1 because the units cancel out? Can someone tell me what it is before the units cancel?

feel like it's kinda straightforward without the step by step conclusion. I'd be comfortable with something like: "To understand and rederive this formula, let's check the dimensions. The left side, I, is charge per time and on the right hand side, q is charge and nvA has dimensions of T^-1."

it doesn't hurt. it's not that long and may be useful to a reader

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_d. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

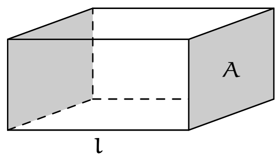

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_d = 6\pi\rho_{fl}\nu v r. \quad (6.6)$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta l r},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will

Comments on page 2

Maybe more time should have been spent on explaining the first method, that is a lot less familiar to me than the check

Is this reading supposed to help us with the pset due tomorrow? I feel like many of the principles overlap.

can you give a real life example of how we will do this? Similar to the example about finding n in a kitchen sink.

I don't get the delta t in the volume formula.

v*delta t gives the width, the distance that the charges have traveled in delta t

This short explanation was helpful to me, maybe this short sentence could be in the text to clarify why there is a delta t in the volume.

basically, we know that the charges are moving at speed v, so over delta(t) they travel the width of the block. also, in terms of dimensional analysis, we want to find a volume here. we need L*L*L. cross sectional area takes care of 2 of those L's. another way to get the 3rd L is v*t.

the delta t is in there because v*deltaT is the width of the rectangular block of volume that we're looking at. As explained earlier, v*deltaT is how far a charge goes in time deltaT.

It does seem strange, just keep in mind that v here is velocity and we need to get rid of the time.

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_d. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

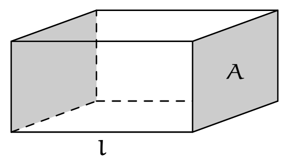

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_d = 6\pi\rho_{fl}\nu v r. \quad (6.6)$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta l r},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will

I know what this looks like since it's been seen in 8.02, but for people who don't remember or others who haven't seen this argument before, a pictoral representation would be extremely helpful here.

Yes, I would also like to see a picture here. But I like how you included this section. I feel like it explains where you pulled the equation for current out of.

i really liked this alternative way of analysis. don't need a picture personally.

This doesn't really say anything new to me that the previous verification didn't. Perhaps you could combine the two.

I disagree, I like the having a separate section. It lets there be a clearer delineation of the sections.

Yeah, this made things very clear to me, I liked this explanation a lot.

I don't think a picture is necessarily mandatory here, just because the algebra is explained pretty thoroughly. However, I guess having the picture doesn't hurt, although it might weaken your argument for "easy cases."

Yeah, the picture already presented was what I just used to think of charge flowing through a block

I'd be helped by a picture. I don't really know much physics.

would it just be a picture of an ampere's loop? don't think thats very necessary.

I think they mean a picture of the block of charges passing through a cross-sectional area A in time delta_t.

I personally think the algebra is sufficient; it's very clear.

I really this this explanation is perfect. Please don't change it.

I remember this now

I really liked this second check. I feel that I'm becoming more familiar with manipulating units.

is this the same for any material? I am assuming yes it is just strange to think about sea water at the conductive material

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_d. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

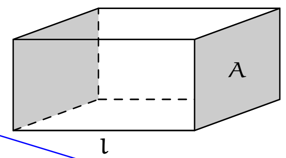

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_d = 6\pi\rho_{fl}\nu v r. \tag{6.6}$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta l r},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will

Specify electrostatic force? + bring up that this is like balancing gravity and drag forces for masses in fluid?

This Fq still confuses me.

is this the speed of electrons? i thought they always went the spped of light?

What exactly is the drag force, and what is causing it?

Just checking, but is the drift speed the same as the charge speed from the preceding paragraph?

I think the drift speed is the average speed, it should be the same as the one in the preceding paragraph. Electrons are the charge carriers in metals and they follow an erratic path, bouncing from atom to atom, but generally drifting in the direction of the electric field. (wiki)

I was wondering about this too. i think they are the same. Drift, in EE, usually refers to something caused by and electric field which is what was considered in the previous paragraph.

This example has gotten kind of 8.02 intensive - I have lost where the simplification is.

Yeah, since some of the paragraphs required some re-reading for me, the problem turned out to be more complex than it actually is.

I agree, I don't know if it's just me but I don't know if I can think of a different method of doing this problem.

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_d. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

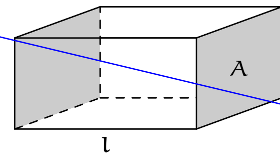

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_d = 6\pi\rho_{fl}\nu v r. \tag{6.6}$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta l r},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will

the divide and conquer tree?

I think so. I don't think there was a point to put the tree at the beginning of the section when we have no clue what it is for.

Perhaps you should show partial trees all the way down, so we can see the tree as you build it up, like you did in the section about trees?

i actually really like the tree at the starting. we've already learned divide and conquer so we can quickly learn where we're going so that the explanations follow suit. I think "preceeding map" should definitely be more specific like "the divide and conquer tree"

Yeah, I don't know if these have been called maps before now.

i agree with how it is now. I took some time and got familiar with where the section was heading and I haven't had to stop and go back or re-read anything

I'm not sure exactly why, but the word "nodes" confused me for a few seconds. Maybe because it is referring to the tree on the previous page? Maybe because I was thinking about circuits and current nodes? It made sense after reading the next sentence and the 3 remaining quantities.

I kind of agree. I actually like the word nodes to describe points on the tree, but either my memory is hazy or you haven't really used node in reference to the tree diagrams in the readings very much. But maybe you should go back and add it in, because it fits with diagrams from comp. sci. and other things.

Ok this answers my earlier question, we still need to calculate n

It's confusing to group these since they are at different parts of the tree. Why not finish out to the leaves and then calculate back up to rho last? After "the preceding map" it might make more sense to jump to "To find v...".

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_d. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

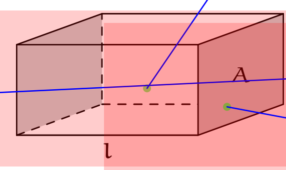

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_d = 6\pi\rho_{fl}\nu v r. \qquad (6.6)$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta l r},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will

This figure doesn't tell me much about rho and R.

I would say this is more confusing than helpful, the way it is currently labeled

It just shows that rho is R times the factor of A/l, which is represented by the box in the figure, but I'm not sure if it really adds any insight that the equation doesn't already do. I guess it just helps to visualize it as well.

how does it show that, exactly?

This figure does little to illustrate anything; I can assume, a bit, what is trying to say to me (in terms of the relationship), but I think just taking it out and taking out the sentence about the figure, or labeling the figure much better, would be better.

I don't see p or R in the figure. How does the figure illustrate their relationship?

Maybe dumb question, but how would one physically measure a block of water and the charge going through it like this?

I don't get were this comes from?

Yeah, I'm surprised there wasn't an explanation given the earlier intro about the different between resistivity and resistance.

This isn't a particularly analytical explanation, but I like to think about it like this: Rewrite it in terms of R= l*rho/A. If there is a large area, it is easier for ions/fluid/whatever to flow through, so resistance based on that geometry would be lower. If the length is long (say, you have a long pipe) it will make flow more difficult than over a very short distance, so resistance should increase.

This is a really helpful way of thinking about it and I feel like it should probably be included in the text. Thanks!

I'm sure we all learned this in 8.02, but short explanations can't hurt.

Why are we multiplying by A? do the other constants above turn into R and we're still left with A?

I agree this relation should have been explained earlier when you explained the tree diagram.

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_d. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

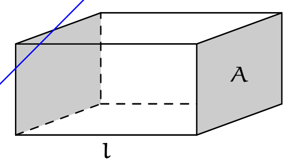

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_d = 6\pi\rho_{fl}\nu vr. \quad (6.6)$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta lr},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will

Why are we using stokes drag here?

We're near the low Reynolds number limit.

this is because ions are small right?

Yes, Re is small because the ions small, the resulting velocity is small, and the viscosity is moderately large.

Good practice is to check this type of assumption after finding a solution.

Thanks for the explanation. I've struggled all year with the "drag" stuff. Bring back the UNIX!

Couldn't disagree with the previous statement more..

When I read this, I wondered this as well.

How many types of drag are there?

Numerically Stokes drag is different than regular air drag, but why is this the case?

I did not find the explanation up to this point clear enough and have to read it again now :(

So we model the charges as solid spheres moving through the water?

It seems that way. It's probably not a bad model. It'd be nice to have a little justification...

I think we're comparing it to the marble in class, which is also a sphere, so modeling the ions as spheres makes it easier.

At least at some minimum length scale, ions appear like points or small spheres, and the only parameter that matters is radial distance from the center.

That's sort of rule 1 for physics: when in doubt, assume your system has the maximum number of symmetries possible, which in this case means assume ions are spheres.

how were we supposed to know this existed?

As a second check, watch a cross-section of the block for a time Δt and calculate how much charge flows during that time. The charges move at speed v, so all charges in a rectangular block of width $v\Delta t$ and area A cross the cross-section. This rectangular block has volume $vA\Delta t$. The ion concentration is n, so the block contains $nvA\Delta t$ charges. If each ion has charge q, then the total charge on the ions is $Q = qnvA\Delta t$. It took a time Δt for this charge to make its journey, so the current is, once again, $I = Q/\Delta t = qnvA$.

The drift speed v depends on the applied force F_q and on the drag force F_d. The ion adjusts its speed until the drag force matches the applied force. The result of this subdividing is the preceding map.

Now let's find expressions for the unknown nodes. Only three remain: ρ, v, and n. The figure illustrates the relation between ρ and R:

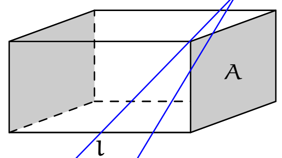

$$\rho = \frac{RA}{l}.$$

To find v, we balance the drag and electrical forces. The applied force is $F_q = qE$, where q is the ion charge and E is the electric field. The electric field produced by the voltage V is $E = V/l$, where l is the length of the block, so

$$F_q = \frac{qV}{l}.$$

This expression contains no unknown quantities, so it does not need further subdivision.

The drag is Stokes drag:

$$F_d = 6\pi\rho_{fl}\nu v r. \tag{6.6}$$

Equating this force to the applied force gives the terminal velocity v in terms of known quantities:

$$v \sim \frac{qV}{6\pi\eta l r},$$

where r is the radius of the ion.

The number density n is the third and final unknown. However, let's estimate it after getting a symbolic result for σ. (This symbolic result will

I really like this method of finding the speed but I think I could have been more simply detailed prior to the method. You mention that we balance the forces, but why?

so what is actually going on?

i'm confused as to where all these (undefined) variables are coming from.

V and l are defined a few paragraphs above, r is defined here, and q is charge. Eta is the dynamic viscosity, and I agree that that switch is unexplained, though it is useful to wrap up the density * nu into one parameter (eta).

Some of the variables tend to be standard when talking about certain types of problems: I will always be current, V voltage, q charge, and so on.

Was wondering about the eta. My question then becomes why is it proportional if they have to equal to balance each other? Are there other forces involved?

The missing step is that Fd = F from the electric field. Then you rearrange terms to get the v from the Fd by itself on one side of the equation. Try writing it out in a few lines, you'll see.

I think that was a rho*nu disappearing into an eta? I guess this is now the dynamic viscosity.

yeah. eta = mu = dynamic viscosity (engineers use mu, scientists use eta... generally).

Symbolic result?

I think he means in terms of variables, instead of just trying to estimate it from the cuff using numbers.

contain n.) To find σ climb the solution tree. First, find the current in terms of the terminal velocity:

$$I = qnvA \sim \frac{q^2nAV}{6\pi\eta lr}.$$

Use the current to find the resistance:

$$R \sim \frac{V}{I} \sim \frac{6\pi\eta lr}{q^2nA}.$$

The voltage V has vanished, which is encouraging: In most circuits, the conductivity (and resistance) is independent of voltage. Use the resistance to find the resistivity:

$$\rho = R\frac{A}{l} \sim \frac{6\pi\eta r}{q^2n}.$$

The expressions simplify as we climb the tree. For example, the geometric parameters l and A have vanished. Their disappearance is encouraging – the purpose of evaluating resistivity rather than resistance is that resistivity is independent of geometry.

Now use resistivity to find conductivity:

$$\sigma = \frac{1}{\rho} \sim \frac{q^2n}{6\pi\eta r}.$$

Here q is the electron charge e or its negative, depending on whether a sodium or a chloride ion is the charge carrier. Thus,

$$\sigma = \frac{1}{\rho} \sim \frac{e^2n}{6\pi\eta r}.$$

To find σ still requires the ion concentration n, which we can find from the concentration of salt in seawater. To do so, try a kitchen-sink experiment. Add table salt to a glass of water until it tastes as salty as seawater. I just tried it. In a glass of water, I found that one teaspoon of salt tastes like drinking seawater. A glass of water may have a volume of $0.3\,\ell$ or a mass of $300\,\text{g}$. A flat teaspoon of salt has a volume of about $5\,\text{m}\ell$. Why $5\,\text{m}\ell$? A teaspoon is about $4\,\text{cm}$ long by $2\,\text{cm}$ wide by $1\,\text{cm}$ thick at its deepest point; let's assume $0.5\,\text{cm}$ on average. Its volume is therefore

$$\text{teaspoon} \sim 4\,\text{cm} \times 2\,\text{cm} \times 0.5\,\text{cm} \sim 4\,\text{cm}^3.$$

It would be nice if the solution tree was reproduced here

That's true, or maybe just the relevant part of it.

I'm not sure I would know to go back through the solution tree and find sigma in order to determine n. Should something be a clue to this or simply the fact that we cannot find it another way?

maybe it would be better to just call it "the tree" instead of "the solution tree"?

Well, that would take up a lot of space, but I definitely agree that some of it should be reproduced for ease ofa ccess.

shouldn't this be = ?

Maybe at some point we decided we were just estimating, and so it might be off by a constant somewhere

But there are no constants in this first part of the equation. R = V / I. Maybe its irrelevant to place an = (since we are only concerned with twiddles), but why would the following equations also use an = and not this one?

Yes.

In what circuits are they dependent on voltage?

I believe the resistance of diodes depend on voltage.

MOSFETS

there are many voltage controlled resistors

Maybe show the step in which it vanishes?

I like the way this is stepped out; it's a good progression and makes it easy to follow. I think the derivations in this section are clear and are important. I'm not lost at all, so I would keep the derivation part as is.

I agree. Sometimes long derivations are easy to get lost in and wonder if you are following correctly. This makes it easy to know that you are on the right path without having to actually work out the derivation on paper

contain n.) To find σ climb the solution tree. First, find the current in terms of the terminal velocity:

$$I = qnvA \sim \frac{q^2nAV}{6\pi\eta lr}.$$

Use the current to find the resistance:

$$R \sim \frac{V}{I} \sim \frac{6\pi\eta lr}{q^2nA}.$$

The voltage V has vanished, which is encouraging: In most circuits, the conductivity (and resistance) is independent of voltage. Use the resistance to find the resistivity:

$$\rho = R\frac{A}{l} \sim \frac{6\pi\eta r}{q^2n}.$$

The expressions simplify as we climb the tree. For example, the geometric parameters l and A have vanished. Their disappearance is encouraging – the purpose of evaluating resistivity rather than resistance is that resistivity is independent of geometry.

Now use resistivity to find conductivity:

$$\sigma = \frac{1}{\rho} \sim \frac{q^2n}{6\pi\eta r}.$$

Here q is the electron charge e or its negative, depending on whether a sodium or a chloride ion is the charge carrier. Thus,

$$\sigma = \frac{1}{\rho} \sim \frac{e^2n}{6\pi\eta r}.$$

To find σ still requires the ion concentration n, which we can find from the concentration of salt in seawater. To do so, try a kitchen-sink experiment. Add table salt to a glass of water until it tastes as salty as seawater. I just tried it. In a glass of water, I found that one teaspoon of salt tastes like drinking seawater. A glass of water may have a volume of $0.3\,\ell$ or a mass of 300 g. A flat teaspoon of salt has a volume of about $5\,m\ell$. Why $5\,m\ell$? A teaspoon is about 4 cm long by 2 cm wide by 1 cm thick at its deepest point; let's assume 0.5 cm on average. Its volume is therefore

$$\text{teaspoon} \sim 4\,\text{cm} \times 2\,\text{cm} \times 0.5\,\text{cm} \sim 4\,\text{cm}^3.$$

A comment about all this math.... if you could add in some reminders around here about what our original goal is, it would help a lot. All this math gets me a little confused, and I spend so much time figuring out what's going on in the math that I forget why we're doing it.

It just makes me get confused and lose interest when I don't know why we are doing the math.

I think the statements in between the lines do enough to clarify the intent.

Yea I think that if you started adding more "clarifying" between equations it would actually be more confusing because the flow of the equations could be lost.

It could be useful to restate the equations that we are using to find these new equations. Even using reference numbers and saying "from equation 1 and 2" could do the trick.

Do you need this word here? I feel you can omit it without any loss of information (pun intended).

It's an important point, though. It implies his next point, that resistivity is independent of geometry. It might be nice to note that, therefore, resistivity is an intensive property (since the terms were brought up before).

you should have written this above as an explanation to the "intensive" property description.

Or just make a note here bringing it back to the earlier point, considering we hadn't derived this yet

I thought this was clearly stated when resistance and resistivity were introduced in the beginning of the section.

I think having a tree in the margin that gets filled in as these values are found would help. The benefit of the tree is lost and confused when the solutions are presented linearly and don't iterate back to the nodes of the tree.

Agreed. I think seeing a tree build up is one of the most useful things - watching them go up in class was usually an epiphany moment to me.

contain n.) To find σ climb the solution tree. First, find the current in terms of the terminal velocity:

$$I = qnvA \sim \frac{q^2 nAV}{6\pi\eta l r}.$$

Use the current to find the resistance:

$$R \sim \frac{V}{I} \sim \frac{6\pi\eta l r}{q^2 nA}.$$

The voltage V has vanished, which is encouraging: In most circuits, the conductivity (and resistance) is independent of voltage. Use the resistance to find the resistivity:

$$\rho = R\frac{A}{l} \sim \frac{6\pi\eta r}{q^2 n}.$$

The expressions simplify as we climb the tree. For example, the geometric parameters l and A have vanished. Their disappearance is encouraging – the purpose of evaluating resistivity rather than resistance is that resistivity is independent of geometry.

Now use resistivity to find conductivity:

$$\sigma = \frac{1}{\rho} \sim \frac{q^2 n}{6\pi\eta r}.$$

Here q is the electron charge e or its negative, depending on whether a sodium or a chloride ion is the charge carrier. Thus,

$$\sigma = \frac{1}{\rho} \sim \frac{e^2 n}{6\pi\eta r}.$$

To find σ still requires the ion concentration n, which we can find from the concentration of salt in seawater. To do so, try a kitchen-sink experiment. Add table salt to a glass of water until it tastes as salty as seawater. I just tried it. In a glass of water, I found that one teaspoon of salt tastes like drinking seawater. A glass of water may have a volume of $0.3\,\ell$ or a mass of 300 g. A flat teaspoon of salt has a volume of about $5\,m\ell$. Why $5\,m\ell$? A teaspoon is about 4 cm long by 2 cm wide by 1 cm thick at its deepest point; let's assume 0.5 cm on average. Its volume is therefore

$$\text{teaspoon} \sim 4\,\text{cm} \times 2\,\text{cm} \times 0.5\,\text{cm} \sim 4\,\text{cm}^3.$$

This makes sense here, maybe explain above that you will use resistivity to find the conductivity

I agree. After reading this part, the intro makes more sense to me when I reread it.

So I looked ahead and saw that finding conductivity is going to take 2.5 pages. It would be great it you could state initially what form we need it in because I got lost in the numbers.

if e is going to be squared, then its sign doesn't matter, does it?

haha, good point. I think that might just have been for explanation of the physical reality

why do we need to do an experiment to calculate concentration?

Do you know the concentration of salt ions in sea water off the top of your head?

is this really necessary?? It's not like this is the type of approximation we can do while taking an exam

I don't think this would be on the exam....and its to demonstrate that it is possible to do simple experiments sometimes to find numbers we want to know.

Yeah, I think this is actually really helpful to see this because too often we jump to the conclusion that we don't know something when there are actually practical ways of guessing.

I don't think the point of this class is to learn how to take an exam.

while I don't think exams are the main focus of this class...I do think this approximation is a bit unnecessary. is it plausible to assume that seawater is saturated with salt?

This sentence is a little weird here. Don't really think it needs to be here.

not necessary but still made me LOL

I think the adds an element of comedy into the text, which is nice. If there are strong objections to having it as a separate sentence, perhaps you could put it in parentheses at the end of the previous sentence and it would flow nicely.

Learning by doing!

contain n.) To find σ climb the solution tree. First, find the current in terms of the terminal velocity:

$$I = qnvA \sim \frac{q^2nAV}{6\pi\eta lr}.$$

Use the current to find the resistance:

$$R \sim \frac{V}{I} \sim \frac{6\pi\eta lr}{q^2nA}.$$

The voltage V has vanished, which is encouraging: In most circuits, the conductivity (and resistance) is independent of voltage. Use the resistance to find the resistivity:

$$\rho = R\frac{A}{l} \sim \frac{6\pi\eta r}{q^2n}.$$

The expressions simplify as we climb the tree. For example, the geometric parameters l and A have vanished. Their disappearance is encouraging – the purpose of evaluating resistivity rather than resistance is that resistivity is independent of geometry.

Now use resistivity to find conductivity:

$$\sigma = \frac{1}{\rho} \sim \frac{q^2n}{6\pi\eta r}.$$

Here q is the electron charge e or its negative, depending on whether a sodium or a chloride ion is the charge carrier. Thus,

$$\sigma = \frac{1}{\rho} \sim \frac{e^2n}{6\pi\eta r}.$$

To find σ still requires the ion concentration n, which we can find from the concentration of salt in seawater. To do so, try a kitchen-sink experiment. Add table salt to a glass of water until it tastes as salty as seawater. I just tried it. In a glass of water, I found that one teaspoon of salt tastes like drinking seawater. A glass of water may have a volume of 0.3 ℓ or a mass of 300 g. A flat teaspoon of salt has a volume of about 5 mℓ. Why 5 mℓ? A teaspoon is about 4 cm long by 2 cm wide by 1 cm thick at its deepest point; let's assume 0.5 cm on average. Its volume is therefore

$$\text{teaspoon} \sim 4\,\text{cm} \times 2\,\text{cm} \times 0.5\,\text{cm} \sim 4\,\text{cm}^3.$$

you may not want to recommend trying this. drinking salt water can cause you to throw up.

A sip of it won't. He's not telling you to chug the whole glass. Besides. You can spit it out after you taste it.

I'm not sure I've tasted seawater enough times to know what it tastes like....

Class field trip to the beach!

Yeah, and I'm not sure that taste alone is enough to actually call it close to seawater.

haha I liked this!

great experimental method!

Why didn't we solve for n at the same time we were solving for other variables, such as rho, earlier in the problem? This seems out of order to me.

why do you need to describe this here? we're doing things far more complicated than volume estimates.

Also, my teaspoons at home say 5mL right on them.

while I need no further proof that I learn at a slower pace than the majority of my classmates, I like it when you break it down into small steps. I'm still baffled when you state facts out of what otherwise seems like omniscience.

it also helps people stay focused on what he is trying to prove instead of getting lost on the fact that a teaspoon is 5mL

I know this is simple, but it is rather nice to have a call-back to what we did at the very beginning!

Sometimes I just sit back and enjoy some of the little, even silly, estimations we do in this class.

The density of salt is maybe twice the density of water, so a flat teaspoon has a mass of roughly 10 g. The mass fraction of salt in seawater is, in this experiment, roughly 1/30. The true value is remarkably close: 0.035. A mole of salt, which provides two charges per NaCl 'molecule', has a mass of 60 g, so

$$n \sim \frac{1}{30} \times \underbrace{1\,\mathrm{g\,cm^{-3}}}_{\rho_{\mathrm{water}}} \times \frac{2\,\mathrm{charges}}{\mathrm{molecule}} \times \frac{6\cdot 10^{23}\,\mathrm{molecules\,mole^{-1}}}{60\,\mathrm{g\,mole^{-1}}}$$
$$\sim 7\cdot 10^{20}\,\mathrm{charges\,cm^{-3}}.$$

With n evaluated, the only remaining mysteries in the conductivity

$$\sigma = \frac{1}{\rho} \sim \frac{q^2 n}{6\pi\eta r}$$

are the ion radius r and the dynamic viscosity η.

Do the easy part first. The dynamic viscosity is

$$\eta = \rho_{\mathrm{water}}\nu \sim 10^3\,\mathrm{kg\,m^{-3}} \times 10^{-6}\,\mathrm{m^2\,s^{-1}} = 10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}.$$

Here I switched to SI (mks) units. Although most calculations are easier in cgs units than in SI units, the one exception is electromagnetism, which is represented by the e^2 in the conductivity. Electromagnetism is conceptually easier in cgs units – which needs no factors of μ_0 or $4\pi\epsilon_0$, for example – than in SI units. However, the cgs unit of charge, the electrostatic unit, is unfamiliar. So, for numerical calculations in electromagnetism, use SI units.

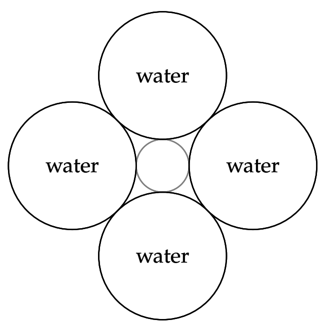

The final quantity required is the ion radius. A positive ion (sodium) attracts an oxygen end of a water molecule; a negative ion (chloride) attracts the hydrogen end of a water molecule. Either way, the ion, being charged, is surrounded by one or maybe more layers of water molecules. As it moves, it drags some of this baggage with it. So rather than use the bare ion radius you should use a larger radius to include this shell. But how thick is the shell? As a guess, assume that the shell includes one layer of water molecules, each with a radius of 1.5 Å. So for the ion plus shell, $r \sim 2$ Å.

Where does this come from? Just an approximation?

I agree without looking that up, i would not have guessed that

most rocks are 2-4 times as dense as water. it's a fair guess

well, from the beginning of the class, water is about 1 and rock is about 3 (g/cm^3), and halite (a rock, or table salt) is probably on the less dense side.

it comes from talking to the gut

talking to the gut?

A reference to a comment in an earlier lecture

It's a reference to the "technique" of getting enough intuition about problems to just pull these sorts of numbers out.

I understand how maybe we should be able to use our gut for this one, but my gut leads me in the wrong direction. Just imagining holding two containers, one with water and one with salt, I want to say the water would be heavier but apparently that's wrong.

thats why you must train your gut by doing several examples and problems and gathering information from around you. the more you do this and talk to your gut, the more accurate it becomes

It's been a while since we've seen something this off-the-cuff and filled with daily experience. I like it.

I find it amazing how these results end up so close to actual values

I think it's also amazing how we can look at a value, like volume of a teaspoon, and have no idea off the top of our heads. then when we do a little bit of applying a known situation and some of our estimation, and we get a feasible answer. also, it's amazing how the estimation of the teaspoon as rectangular works so well.

Why 1/30? It seems like most predictions used in this class are often nicer than this.

This was an extremely clever way to get this value. I would have never thought of making my own seawater to find the mass fraction of salt in it. Cool stuff.

The density of salt is maybe twice the density of water, so a flat teaspoon has a mass of roughly 10 g. The mass fraction of salt in seawater is, in this experiment, roughly 1/30. The true value is remarkably close: 0.035. A mole of salt, which provides two charges per NaCl 'molecule', has a mass of 60 g, so

$$n \sim \frac{1}{30} \times \underbrace{1\,\mathrm{g\,cm^{-3}}}_{\rho_{\mathrm{water}}} \times \frac{2\,\mathrm{charges}}{\mathrm{molecule}} \times \frac{6\cdot 10^{23}\,\mathrm{molecules\,mole^{-1}}}{60\,\mathrm{g\,mole^{-1}}}$$
$$\sim 7\cdot 10^{20}\,\mathrm{charges\,cm^{-3}}.$$

With n evaluated, the only remaining mysteries in the conductivity

$$\sigma = \frac{1}{\rho} \sim \frac{q^2 n}{6\pi\eta r}$$

are the ion radius r and the dynamic viscosity η.

Do the easy part first. The dynamic viscosity is

$$\eta = \rho_{\mathrm{water}}\nu \sim 10^3\,\mathrm{kg\,m^{-3}} \times 10^{-6}\,\mathrm{m^2\,s^{-1}} = 10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}.$$

Here I switched to SI (mks) units. Although most calculations are easier in cgs units than in SI units, the one exception is electromagnetism, which is represented by the e^2 in the conductivity. Electromagnetism is conceptually easier in cgs units – which needs no factors of μ_0 or $4\pi\epsilon_0$, for example – than in SI units. However, the cgs unit of charge, the electrostatic unit, is unfamiliar. So, for numerical calculations in electromagnetism, use SI units.

The final quantity required is the ion radius. A positive ion (sodium) attracts an oxygen end of a water molecule; a negative ion (chloride) attracts the hydrogen end of a water molecule. Either way, the ion, being charged, is surrounded by one or maybe more layers of water molecules. As it moves, it drags some of this baggage with it. So rather than use the bare ion radius you should use a larger radius to include this shell. But how thick is the shell? As a guess, assume that the shell includes one layer of water molecules, each with a radius of 1.5 Å. So for the ion plus shell, $r \sim 2$ Å.

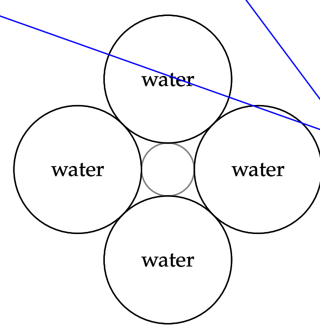

This part is the hardest for me to figure out (chemistry was never my favorite). Where does this number come from?

A single NaCl pair will split up in water into an Na+ and a Cl- (the Na donates one electron to the Cl), each of which has a charge of +e or -e (the charge of one electron). Thus, we now have two charges which can move around and conduct electricity.

why is molecule in 's?

In solution and in crystals they are separate ions, so even though they come 1:1, they aren't considered a molecule.

I haven't seen chemistry in a while and i can't exactly remember how much a mole is..

a mole is 6*10^23. the reason we're using moles here is because we know molar mass, so we can divide the two

OK so this was very well explained. I cringed at the thought of two pages to explain conductivity, but it doesn't seem too bad.

when did this appear?

this variable should have been defined earlier...

I'm not familiar with this unit, whats mks stand for?

i hadn't even really realized that you'd changed units until you mentioned it. i'm not sure you need to explicitly state and explain why you choose to switch units.

i think it's a very important thing to keep track of. half of this entire course is making sure that the units match up, and this is an example of that.

yeah I agree, its really important to realize what units you're working in so this sentence is really helpful

It'd be helpful if you mentioned what cgs units are because I have no clue. Is it some common knowledge I am missing?

It's a convention for units referring to centimeters, grams, seconds.

interesting point, I would have just tried to continue using cgs, and avoided the confusion of having to switch back and forth.

The density of salt is maybe twice the density of water, so a flat teaspoon has a mass of roughly 10 g. The mass fraction of salt in seawater is, in this experiment, roughly 1/30. The true value is remarkably close: 0.035. A mole of salt, which provides two charges per NaCl 'molecule', has a mass of 60 g, so

$$n \sim \frac{1}{30} \times \underbrace{1\,\text{g}\,\text{cm}^{-3}}_{\rho_{\text{water}}} \times \frac{2\,\text{charges}}{\text{molecule}} \times \frac{6 \cdot 10^{23}\,\text{molecules mole}^{-1}}{60\,\text{g mole}^{-1}}$$
$$\sim 7 \cdot 10^{20}\,\text{charges}\,\text{cm}^{-3}.$$

With n evaluated, the only remaining mysteries in the conductivity

$$\sigma = \frac{1}{\rho} \sim \frac{q^2 n}{6\pi\eta r}$$

are the ion radius r and the dynamic viscosity η.

Do the easy part first. The dynamic viscosity is

$$\eta = \rho_{\text{water}}\nu \sim 10^3\,\text{kg}\,\text{m}^{-3} \times 10^{-6}\,\text{m}^2\,\text{s}^{-1} = 10^{-3}\,\text{kg}\,\text{m}^{-1}\,\text{s}^{-1}.$$

Here I switched to SI (mks) units. Although most calculations are easier in cgs units than in SI units, the one exception is electromagnetism, which is represented by the e^2 in the conductivity. Electromagnetism is conceptually easier in cgs units – which needs no factors of μ_0 or $4\pi\epsilon_0$, for example – than in SI units. However, the cgs unit of charge, the electrostatic unit, is unfamiliar. So, for numerical calculations in electromagnetism, use SI units.

The final quantity required is the ion radius. A positive ion (sodium) attracts an oxygen end of a water molecule; a negative ion (chloride) attracts the hydrogen end of a water molecule. Either way, the ion, being charged, is surrounded by one or maybe more layers of water molecules. As it moves, it drags some of this baggage with it. So rather than use the bare ion radius you should use a larger radius to include this shell. But how thick is the shell? As a guess, assume that the shell includes one layer of water molecules, each with a radius of 1.5 Å. So for the ion plus shell, $r \sim 2$ Å.

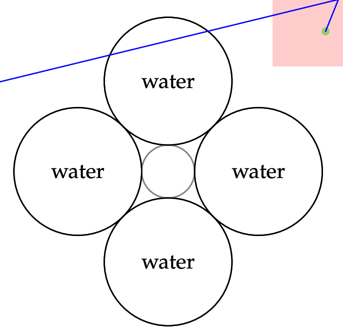

what are cgs units? i'm assuming the English system...

No, it's cm/g/sec vs. m/kg/s for SI. It's the EM units that are more different.

As an 8.022 buff, I disagree. I actually found cgs units better (for the reasons you mentioned). If you need to, you can convert from cgs to mks at the end.

I like cgs also, but conversions are tricky because the units are actually fundamentally different. It's not like converting feet to meters, it's almost like converting feet to kilograms. That's not the best analogy, since there is a conversion, but it's more like values in cgs correspond to values in SI.

Earlier you said the last unknown was n. I think this explanation could use some restructuring for consistency and to better reflect the tree.

nice illustration here

I agree as well, simple but effective!

This definitely makes sense, but perhaps you could help us visualize it with a picture.

This is a pretty hard problem. I am encouraged by the attempts to simplify solvation...

I agree, I kind of assumed this but any type of approximation skills I have are greatly undermined when I have to give a value for the radius of an ion.

Wouldn't there be some repulsion to keep it away? ie. the hydrogen will move away from the positive sodium and keep the molecules not to far. Or will the molecules constantly be moving trying to find the unit they are attracted to?

The density of salt is maybe twice the density of water, so a flat teaspoon has a mass of roughly 10 g. The mass fraction of salt in seawater is, in this experiment, roughly 1/30. The true value is remarkably close: 0.035. A mole of salt, which provides two charges per NaCl 'molecule', has a mass of 60 g, so

$$n \sim \frac{1}{30} \times \underbrace{1\,\mathrm{g\,cm^{-3}}}_{\rho_{\mathrm{water}}} \times \frac{2\,\mathrm{charges}}{\mathrm{molecule}} \times \frac{6 \cdot 10^{23}\,\mathrm{molecules\,mole^{-1}}}{60\,\mathrm{g\,mole^{-1}}}$$
$$\sim 7 \cdot 10^{20}\,\mathrm{charges\,cm^{-3}}.$$

With n evaluated, the only remaining mysteries in the conductivity

$$\sigma = \frac{1}{\rho} \sim \frac{q^2 n}{6\pi\eta r}$$

are the ion radius r and the dynamic viscosity η.

Do the easy part first. The dynamic viscosity is

$$\eta = \rho_{\mathrm{water}}\nu \sim 10^3\,\mathrm{kg\,m^{-3}} \times 10^{-6}\,\mathrm{m^2\,s^{-1}} = 10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}.$$

Here I switched to SI (mks) units. Although most calculations are easier in cgs units than in SI units, the one exception is electromagnetism, which is represented by the e^2 in the conductivity. Electromagnetism is conceptually easier in cgs units – which needs no factors of μ_0 or $4\pi\epsilon_0$, for example – than in SI units. However, the cgs unit of charge, the electrostatic unit, is unfamiliar. So, for numerical calculations in electromagnetism, use SI units.

The final quantity required is the ion radius. A positive ion (sodium) attracts an oxygen end of a water molecule; a negative ion (chloride) attracts the hydrogen end of a water molecule. Either way, the ion, being charged, is surrounded by one or maybe more layers of water molecules. As it moves, it drags some of this baggage with it. So rather than use the bare ion radius you should use a larger radius to include this shell. But how thick is the shell? As a guess, assume that the shell includes one layer of water molecules, each with a radius of 1.5 Å. So for the ion plus shell, $r \sim 2$ Å.

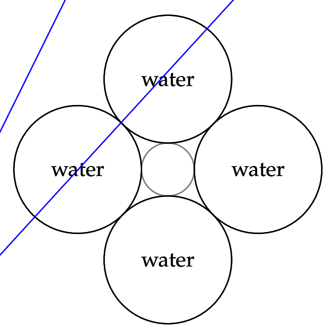

This is so involved- I definitely would not have thought of doing this to calculate the ion radius

I agree. Although I think this is cool, this problem transcends basic approximation to become more of a technical science problem

Yeah there is a lot more 8.01,8.02 and 5.111 in this problem than there should be for this to be considered a basic approximation problem.

Yup, I never would have thought of it either. On the other hand, completely ignoring this fact would have probably led me to guess an ion radius of about .5 or 1 Angstrom, which is only off by a factor of 2. So the danger from missing this fact is pretty low.

Wow very true...I guess that's why its our homework to read and learn about it.

Agreed, so makes me wonder, would we actually include this if we were calculating this? i understand why it's included in the text for completeness, but perhaps it should be mentioned, after arriving at the answer, that perhaps intricacies like this could be skipped over next time?

We found before that a hydrogen atom has a diameter of 1A. How could we extrapolate to this number?

This confused me also. I think earlier we did the hydrogen atom (not ion). I think this is one of those numbers you may have to look up.

But what is the deciding factor in looking up a value? Looking up the volume of a teaspoon would have been just as easy...

You could model water as a space-filling model (little Mickey Mouse with hydrogen ears) or ball + stick model, so its radius would be something a little bigger than 1A. http://www.sacredbalance.com/web/flashplayer.html?id=h2o

I agree. You know that the hydrogen atoms sit on the oxygen, not straight across but at an angle. If you really remembered your chemistry you would know that angle, but even without it you can assume if each atom is 1A across, then the total has to be less than 3, so maybe a result in the 1.5-2A range is reasonable.

With these numbers, the conductivity becomes:

$$\sigma \sim \frac{\overbrace{(1.6\cdot 10^{-19}\,\mathrm{C})^2}^{e^2} \times \overbrace{7\cdot 10^{26}\,\mathrm{m}^{-3}}^{n}}{\underbrace{6\times 3}_{6\pi} \times \underbrace{10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}}_{\eta} \times \underbrace{2\cdot 10^{-10}\,\mathrm{m}}_{r}}.$$

You can do the computation mentally. First count the powers of ten and then worrying about the small factors. Then count the top and bottom contributions separately. The top contributes −12 powers of 10: −38 from e^2 and +26 from n. The bottom contributes −13 powers of 10: −3 from η and −10 from r. The division produces one power of 10.

Now account for the remaining small factors:

$$\frac{1.6^2 \times 7}{6\times 3\times 2}.$$

Slightly overestimate the answer by pretending that the 1.6^2 on top cancels the 3 on the bottom. Slightly underestimate the answer – and maybe compensate for the overestimate – by pretending that the 7 on top cancels the 6 on the bottom. After these lies, only 1/2 remains. Multiplying it by the sole power of ten gives

$$\sigma \sim 5\,\Omega^{-1}\,\mathrm{m}^{-1}.$$

Using a calculator to do the arithmetic gives $4.977\ldots\,\Omega^{-1}\,\mathrm{m}^{-1}$, which is extremely close to the result from mental calculation.

The estimated resistivity is

$$\rho \sim \sigma^{-1} \sim 0.2\,\Omega\,\mathrm{m} = 20\,\Omega\,\mathrm{cm},$$

where we converted to the conventional although not fully SI units of Ω cm. A typical experimental value for seawater at $T = 15\,^\circ\mathrm{C}$ is $23.3\,\Omega\,\mathrm{cm}$ (from [15, p. **14**-15]), absurdly close to the estimate!

Probably the most significant error is the radius of the ion-plus-water combination that is doing the charge transport. Perhaps r should be greater than 2 Å, especially for a sodium ion, which is smaller than chloride; it therefore has a higher electric field at its surface and grabs water molecules more strongly than chloride does. In spite of such uncertainties, the continuum approximation produced accurate results.

This notation is EXTREMELY helpful in understanding the equation- I would love to see this in other equations!

Agreed! This is useful, especially if its 12am and I can't remember exactly what I read a few pages ago

I enjoy these approximations from the beginning.

Yes me too. Can we have more readings like this where we incorporate past concepts? I think it would hone in everything better.

I cannot agree to this enough (even if it's difficult to function do in latex)

yeah i definitely agree too! sometimes after a long derivation, a big equation with numbers is shown and I don't really remember what each number is referring to, but this way of representing the equation makes it very clear!

I think this is a good touch because it relates to thinks we did before. It's always good to see a course building among itself.

This sentence doesn't make sense?

I think it makes sense once you change "worrying" to "worry."

I like that this references back to one of the first lessons!

Agreed. It never hurts to a refreshing lesson every now and then.

I find that I make so many more careless mistakes when I do things in my head instead of writing them down. Does it still count as doing it in your head if you write down one or two numbers to help yourself along as you go?

I think this method goes without saying..

i like that you explain this again here. it reminds us of the simple stuff we learned at the beginning and how it's always useful.

I also find the quick review nice, it's easy to forget with all the different approximation methods we've used in between.

Very true, going back to the first few things we learned keeps us on our feet

With these numbers, the conductivity becomes:

$$\sigma \sim \frac{\overbrace{(1.6\cdot 10^{-19}\,\mathrm{C})^2}^{e^2} \times \overbrace{7\cdot 10^{26}\,\mathrm{m}^{-3}}^{n}}{\underbrace{6\times 3}_{6\pi} \times \underbrace{10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}}_{\eta} \times \underbrace{2\cdot 10^{-10}\,\mathrm{m}}_{r}}.$$

You can do the computation mentally. First count the powers of ten and then worrying about the small factors. Then count the top and bottom contributions separately. The top contributes −12 powers of 10: −38 from e^2 and +26 from n. The bottom contributes −13 powers of 10: −3 from η and −10 from r. The division produces one power of 10.

Now account for the remaining small factors:

$$\frac{1.6^2\times 7}{6\times 3\times 2}.$$

Slightly overestimate the answer by pretending that the 1.6^2 on top cancels the 3 on the bottom. Slightly underestimate the answer – and maybe compensate for the overestimate – by pretending that the 7 on top cancels the 6 on the bottom. After these lies, only 1/2 remains. Multiplying it by the sole power of ten gives

$$\sigma \sim 5\,\Omega^{-1}\,\mathrm{m}^{-1}.$$

Using a calculator to do the arithmetic gives $4.977\ldots\,\Omega^{-1}\,\mathrm{m}^{-1}$, which is extremely close to the result from mental calculation.

The estimated resistivity is

$$\rho \sim \sigma^{-1} \sim 0.2\,\Omega\,\mathrm{m} = 20\,\Omega\,\mathrm{cm},$$

where we converted to the conventional although not fully SI units of Ω cm. A typical experimental value for seawater at $T = 15\,°\mathrm{C}$ is $23.3\,\Omega$ cm (from [15, p. **14**-15]), absurdly close to the estimate!

Probably the most significant error is the radius of the ion-plus-water combination that is doing the charge transport. Perhaps r should be greater than 2 Å, especially for a sodium ion, which is smaller than chloride; it therefore has a higher electric field at its surface and grabs water molecules more strongly than chloride does. In spite of such uncertainties, the continuum approximation produced accurate results.

Where did we define Omega?

I think he means Ohms, which is a measurement of resistance. It may help to state that though.

is there something up with the NB system. i just tried deleting the comment right under this one and it doesn't seem to fully work.

Often times you have to refresh the page to make changes like deleted comments actually register. The NB website doesn't update in realtime, it just loads the most current version of the document when you sign on.

I find it funny after the long argument to use mental math, you resort to a calculator in the end. This is much easier to enter though, so I see the point.

I like how you encouraged mental calculation. I feel like the computing part of our brain is atrophying because of our dependence on machines to do even simple calculations.

Yep - this part is good too since it definitely ties back into the "few" arguments we were making at the beginning of the term.

I completely agree with this I often find myself doing the simplest calculations on a calculator which ends up making more mistakes due to typos and it's good to have a class where we exercise our brains again

what do you mean by this?

ohms and cm are not SI. mhos and meters are.

With these numbers, the conductivity becomes:

$$\sigma \sim \frac{\overbrace{(1.6 \cdot 10^{-19}\,\mathrm{C})^2}^{e^2} \times \overbrace{7 \cdot 10^{26}\,\mathrm{m}^{-3}}^{n}}{\underbrace{6 \times 3}_{6\pi} \times \underbrace{10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}}_{\eta} \times \underbrace{2 \cdot 10^{-10}\,\mathrm{m}}_{r}}.$$

You can do the computation mentally. First count the powers of ten and then worrying about the small factors. Then count the top and bottom contributions separately. The top contributes −12 powers of 10: −38 from e^2 and +26 from n. The bottom contributes −13 powers of 10: −3 from η and −10 from r. The division produces one power of 10.

Now account for the remaining small factors:

$$\frac{1.6^2 \times 7}{6 \times 3 \times 2}.$$

Slightly overestimate the answer by pretending that the 1.6^2 on top cancels the 3 on the bottom. Slightly underestimate the answer – and maybe compensate for the overestimate – by pretending that the 7 on top cancels the 6 on the bottom. After these lies, only 1/2 remains. Multiplying it by the sole power of ten gives

$$\sigma \sim 5\,\Omega^{-1}\,\mathrm{m}^{-1}.$$

Using a calculator to do the arithmetic gives $4.977\ldots\,\Omega^{-1}\,\mathrm{m}^{-1}$, which is extremely close to the result from mental calculation.

The estimated resistivity is

$$\rho \sim \sigma^{-1} \sim 0.2\,\Omega\,\mathrm{m} = 20\,\Omega\,\mathrm{cm},$$

where we converted to the conventional although not fully SI units of Ω cm. A typical experimental value for seawater at $T = 15\,^\circ\mathrm{C}$ is $23.3\,\Omega$ cm (from [15, p. **14-15**]), absurdly close to the estimate!

Probably the most significant error is the radius of the ion-plus-water combination that is doing the charge transport. Perhaps r should be greater than 2 Å, especially for a sodium ion, which is smaller than chloride; it therefore has a higher electric field at its surface and grabs water molecules more strongly than chloride does. In spite of such uncertainties, the continuum approximation produced accurate results.

Why did we not figure temperature in?

Maybe resistivity changes? Or ion density?

I'm not sure. From http://sam.ucsd.edu/sio210/lect_2/lecture_2.html

" Conductivity of sea water depends strongly on temperature, somewhat less strongly on salinity, and very weakly on pressure. If the temperature is measured, then conductivity can be used to determine the salinity. Salinity as computed through conductivity appears to be more closely related to the actual dissolved constituents than is chlorinity, and more independent of salt composition. Therefore temperature must be measured at the same time as conductivity, to remove the temperature effect and obtain salinity. Accuracy of salinity determined from conductivity: 0.001 to 0.004. Precision: 0.001. The accuracy depends on the accuracy of the seawater standard used to calibrate the conductivity based measurement. "

this is crazy how close it is

wow that is impressive

I really like this outline of the possible flaws of the estimation.

I agree, I don't think I'd be able to get through of these steps on my own but I can understand the reasoning because of this.

Since the net charge is the same, the field as a function of distance from the center of the ion should be the same (outside the volume of the ion itself, that is). I would have guessed that the size of the water/ion ball, therefore, wouldn't depend on the ion size (unless the ion were very large).

I'm assuming here that the radius of the water/ion ball is determined by when the electric field drops below some value, and can no longer attract water molecules. Perhaps that assumption is flawed.

With these numbers, the conductivity becomes:

$$\sigma \sim \frac{\overbrace{(1.6\cdot 10^{-19}\,\mathrm{C})^2}^{e^2} \times \overbrace{7\cdot 10^{26}\,\mathrm{m}^{-3}}^{n}}{\underbrace{6\times 3}_{6\pi} \times \underbrace{10^{-3}\,\mathrm{kg\,m^{-1}\,s^{-1}}}_{\eta} \times \underbrace{2\cdot 10^{-10}\,\mathrm{m}}_{r}}.$$

You can do the computation mentally. First count the powers of ten and then worrying about the small factors. Then count the top and bottom contributions separately. The top contributes −12 powers of 10: −38 from e^2 and +26 from n. The bottom contributes −13 powers of 10: −3 from η and −10 from r. The division produces one power of 10.

Now account for the remaining small factors:

$$\frac{1.6^2 \times 7}{6\times 3\times 2}.$$

Slightly overestimate the answer by pretending that the 1.6^2 on top cancels the 3 on the bottom. Slightly underestimate the answer – and maybe compensate for the overestimate – by pretending that the 7 on top cancels the 6 on the bottom. After these lies, only 1/2 remains. Multiplying it by the sole power of ten gives

$$\sigma \sim 5\,\Omega^{-1}\,\mathrm{m}^{-1}.$$

Using a calculator to do the arithmetic gives $4.977\ldots\,\Omega^{-1}\,\mathrm{m}^{-1}$, which is extremely close to the result from mental calculation.

The estimated resistivity is

$$\rho \sim \sigma^{-1} \sim 0.2\,\Omega\,\mathrm{m} = 20\,\Omega\,\mathrm{cm},$$

where we converted to the conventional although not fully SI units of $\Omega\,\mathrm{cm}$. A typical experimental value for seawater at $T = 15\,^\circ\mathrm{C}$ is $23.3\,\Omega\,\mathrm{cm}$ (from [15, p. **14**-15]), absurdly close to the estimate!

Probably the most significant error is the radius of the ion-plus-water combination that is doing the charge transport. Perhaps r should be greater than 2 Å, especially for a sodium ion, which is smaller than chloride; it therefore has a higher electric field at its surface and grabs water molecules more strongly than chloride does. In spite of such uncertainties, the continuum approximation produced accurate results.

cool results! but how can this be considered a "simple case." it just took you 6 pages to cover it, whereas some of the other examples we've seen takes like half a page.

I agree. I'm also wondering, are we still looking at lossy examples here? I'm confused as to what terminology you're using for all of these.

Although this section took me a couple reads, I think it definitely still qualifies as an "easy case." Take a look at the question again, it's pretty complex. But we were still able to find an answer by drawing a tree diagram and using simple cases. Although it took 6 pages, you have to remember that this particular problem required a lot of algebra, which is why it spanned so many pages (plus it was thoroughly explained which was nice).

As I understand it, "simple cases" doesn't mean "doing problems which are simple." Rather, it seems to mean "doing hard problems using the simplest possible setup." In this case, we chose a very simple setup of imagining conduction between two plates of area A, which made all the geometric calculation much easier.

That accuracy is puzzling. At the length scale of a sodium ion, water looks like a collection of spongy boulders more than it looks like a continuum. Yet Stokes drag worked. It works because the important length scale is not the size of water molecules, but rather their mean free path between collisions. Molecules in a liquid are packed to the point of contact, so the mean free path is much shorter than a molecular (or even ionic) radius, especially compared to an ion with its shell of water.

The moral of this example, besides the application of Stokes drag, is to have courage: Approximate first and ask questions later. The approximations might be accurate for reasons that you do not suspect when you start solving a problem. If you agonize over each approximation, you will never start a calculation, and then you will not find out that many approximations would have been fine – if only you had had the courage to make them.

6.3.4 Combining solutions from the two limits

You know know the drag force in two extreme cases, viscous and turbulent drag. The results are repeated here:

$$F_d = \begin{cases} 6\pi\rho_{fl}\nu v r & \text{(viscous)}, \\ \frac{1}{2}c_d\rho_{fl}Av^2 & \text{(turbulent)}. \end{cases}$$

Let's compare and combine them by making the viscous form look like the turbulent form: Multiply by the Reynolds number rv/ν (basing the Reynolds number on radius rather than diameter). Then

$$F_d = \underbrace{\left(\frac{rv/\nu}{\text{Re}}\right)}_{1} \times \underbrace{6\pi\rho_{fl}\nu v r}_{F_d} = \frac{1}{\text{Re}}\,6\pi\rho_{fl}v^2r^2 \qquad \text{(viscous)}.$$

With $A = \pi r^2$,

$$F_d = \frac{6}{\text{Re}}\rho_{fl}v^2A \qquad \text{(viscous)},$$

Since $c_d \equiv F_d/(\frac{1}{2}\rho_{fl}v^2A)$,

$$c_d = \frac{12}{\text{Re}} \qquad \text{(viscous)}.$$

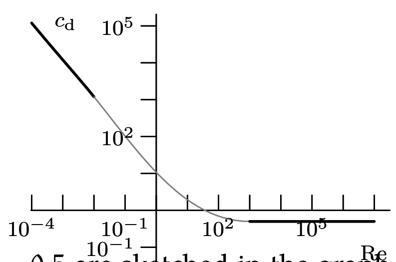

That limit and the the high-speed limit $c_d \sim 0.5$ are sketched in the graph (with a gray interpolation between the limits). Almost all of the experimental data is explained by this graph, except for the drop in c_d near the $\text{Re} \sim 10^6$.

Comments on page 6

I was just about to leave a comment to this effect.

I feel like all the calculations done in these notes turn out eerily accurate- if I were to try and estimate the same values I would probably be way more off.

I wonder what the average percent accuracy is with these estimation (such as, answers differed from actual results by x%)

did we assume stokes drag because of the low reynolds number?

How does short mean free path make it more like a continuum? I have a gut agreement with this. Is it the finer scale, the way moving my hand through sand is more like a continuum than moving through gravel?

Will this idea be repeated for future concepts?

yea, i think that's become very clear over this course!

But what if your approximation wasn't done correctly, but still numerically accurate?

I really like how this tends to work out, we've seen it a couple times now

I like this focus on the ends instead of the means

This seems applicable to life decisions in general haha

I really like this paragraph. This has always been my main problem with approximations in the past. This is something this class has taught me and even if it can get some wrong answers, they eventually lead to correct ones.

That accuracy is puzzling. At the length scale of a sodium ion, water looks like a collection of spongy boulders more than it looks like a continuum. Yet Stokes drag worked. It works because the important length scale is not the size of water molecules, but rather their mean free path between collisions. Molecules in a liquid are packed to the point of contact, so the mean free path is much shorter than a molecular (or even ionic) radius, especially compared to an ion with its shell of water.

The moral of this example, besides the application of Stokes drag, is to have courage: Approximate first and ask questions later. The approximations might be accurate for reasons that you do not suspect when you start solving a problem. If you agonize over each approximation, you will never start a calculation, and then you will not find out that many approximations would have been fine – if only you had had the courage to make them.

6.3.4 Combining solutions from the two limits

You know know the drag force in two extreme cases, viscous and turbulent drag. The results are repeated here:

$$F_d = \begin{cases} 6\pi\rho_{fl}\nu v r & \text{(viscous)},\\ \frac{1}{2}c_d\rho_{fl}Av^2 & \text{(turbulent)}.\end{cases}$$

Let's compare and combine them by making the viscous form look like the turbulent form: Multiply by the Reynolds number rv/ν (basing the Reynolds number on radius rather than diameter). Then

$$F_d = \underbrace{\left(\frac{rv/\nu}{\text{Re}}\right)}_{1} \times \underbrace{6\pi\rho_{fl}\nu v r}_{F_d} = \frac{1}{\text{Re}}\,6\pi\rho_{fl}v^2r^2 \qquad \text{(viscous)}.$$

With $A = \pi r^2$,

$$F_d = \frac{6}{\text{Re}}\rho_{fl}v^2A \qquad \text{(viscous)},$$

Since $c_d \equiv F_d/(\frac{1}{2}\rho_{fl}v^2A)$,

$$c_d = \frac{12}{\text{Re}} \qquad \text{(viscous)}.$$

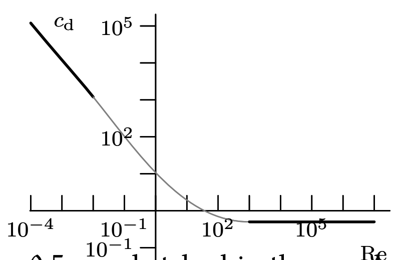

That limit and the the high-speed limit $c_d \sim 0.5$ are sketched in the graph (with a gray interpolation between the limits). Almost all of the experimental data is explained by this graph, except for the drop in c_d near the $\text{Re} \sim 10^6$.

So.... yet again, how is this "easy cases"? Nothing about this really seemed "easy". I was pretty confused throughout. Also, I'm still not sure how this unit is about lossy complexity.

i agree...

I was thinking this the whole time I was reading this. I guess the easy case was stated in the beginning that we are looking at low Reynolds number. But no where in the reading did he expound on the fact that this conclusion came about because of this easy case.

I'm going to be so lost when we get a hard case.

In real life how would you normally go about finding out if you were close to the true answer? It seems like in a lot of situations you might not ever find out...

deep, I can totally relate to times when i didnt even know how to start approximating

This is how I feel about the homework assignments! I feel like each time I surprise myself by being close, I'm more likely to trust my gut instinct.

who would have known that my engineering class was meant to teach me courage?

It is better to have approximated and lost (accuracy) than to have never approximated at all...

Haha it's funny, his word choice reminded me a lot of that love quote too.

Maybe here you could have a quick recap of how our results relate to the Reynolds number, especially before we try to integrate these results in the next section.

Agreed. At least we should check the laminar-regime assumption.

could you use this same reasoning to approximate the resistivity of any material? and I wish you explained resistivity a little better

That accuracy is puzzling. At the length scale of a sodium ion, water looks like a collection of spongy boulders more than it looks like a continuum. Yet Stokes drag worked. It works because the important length scale is not the size of water molecules, but rather their mean free path between collisions. Molecules in a liquid are packed to the point of contact, so the mean free path is much shorter than a molecular (or even ionic) radius, especially compared to an ion with its shell of water.

The moral of this example, besides the application of Stokes drag, is to have courage: Approximate first and ask questions later. The approximations might be accurate for reasons that you do not suspect when you start solving a problem. If you agonize over each approximation, you will never start a calculation, and then you will not find out that many approximations would have been fine – if only you had had the courage to make them.

6.3.4 Combining solutions from the two limits

You know know the drag force in two extreme cases, viscous and turbulent drag. The results are repeated here:

$$F_d = \begin{cases} 6\pi\rho_{fl}\nu v r & \text{(viscous)}, \\ \frac{1}{2}c_d\rho_{fl}Av^2 & \text{(turbulent)}. \end{cases}$$

Let's compare and combine them by making the viscous form look like the turbulent form: Multiply by the Reynolds number rv/ν (basing the Reynolds number on radius rather than diameter). Then

$$F_d = \underbrace{\left(\frac{rv/\nu}{\text{Re}}\right)}_{1} \times \underbrace{6\pi\rho_{fl}\nu v r}_{F_d} = \frac{1}{\text{Re}}\,6\pi\rho_{fl}v^2r^2 \qquad \text{(viscous)}.$$

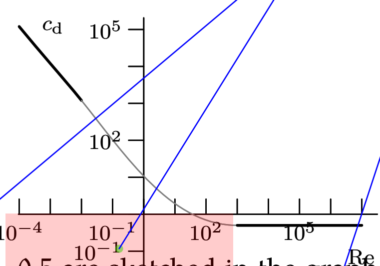

With $A = \pi r^2$,

$$F_d = \frac{6}{\text{Re}}\rho_{fl}v^2A \qquad \text{(viscous)},$$

Since $c_d \equiv F_d/(\frac{1}{2}\rho_{fl}v^2A)$,

$$c_d = \frac{12}{\text{Re}} \qquad \text{(viscous)}.$$

That limit and the the high-speed limit $c_d \sim 0.5$ are sketched in the graph (with a gray interpolation between the limits). Almost all of the experimental data is explained by this graph, except for the drop in c_d near the $\text{Re} \sim 10^6$.

I like this section- to the point and very understandable.

I agree. It is also interesting to see how we are able to come from those two equations to make a coherent plot.

Yup, although very brief, it was really nice at summing everything up and tying it all together.

I also strongly agree. I actually got so caught up in details above that I had lost sight of the point of what we were doing. This section nicely brought everything back together.

Is this supposed to be "now"?

why are we doing this?

just a note, the graph is too close to the writing.

only near 10ˆ6, how about greater than that? why is this an exception

I believe there is a plan afoot to show us this soon.

Did I miss something, or was there no actual calculation of Re?

7
Lumping

7.1 Estimating populations: How many babies? 133
7.2 Bending of light 135
7.3 Quantum mechanics 138
7.4 Sound and electromagnetic radiation 143
7.5 Boundary layers 143

The symmetry chapter (Section 3.1) introduced the principle of invariance: 'When there is change, look for what does not change.' However, when you cannot find any useful but unchanging quantity, you have to make one. As Jean-Luc Picard often says, 'Make it so.'

7.1 Estimating populations: How many babies?

The first example is to estimate the number of babies in the United States. For definiteness, call a child a baby until he or she turns 2 years old. An exact calculation requires the birth dates of every person in the United States. This, or closely similar, information is collected once every decade by the US Census Bureau.

As an approximation to this voluminous data, the Census Bureau [33] publishes the number of people at each age. The data for 1991 is a set of points lying on a wiggly line $N(t)$, where t is age. Then

$$N_{\text{babies}} = \int_0^{2\,\text{yr}} N(t)\,dt. \tag{7.1}$$

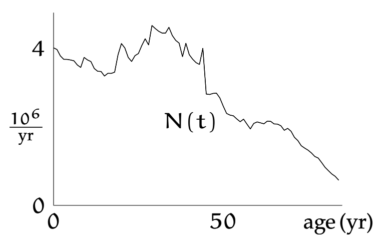

I like how this section starts with a very easy example and then goes to a more complicated one and then gives a more difficult practice problem.

7
Lumping

7.1 Estimating populations: How many babies? 133
7.2 Bending of light 135
7.3 Quantum mechanics 138
7.4 Sound and electromagnetic radiation 143
7.5 Boundary layers 143

The symmetry chapter (Section 3.1) introduced the principle of invariance: 'When there is change, look for what does not change.' However, when you cannot find any useful but unchanging quantity, you have to make one. As Jean-Luc Picard often says, 'Make it so.'

7.1 Estimating populations: How many babies?

The first example is to estimate the number of babies in the United States. For definiteness, call a child a baby until he or she turns 2 years old. An exact calculation requires the birth dates of every person in the United States. This, or closely similar, information is collected once every decade by the US Census Bureau.

As an approximation to this voluminous data, the Census Bureau [33] publishes the number of people at each age. The data for 1991 is a set of points lying on a wiggly line $N(t)$, where t is age. Then

$$N_{babies} = \int_0^{2\,yr} N(t)\,dt. \tag{7.1}$$

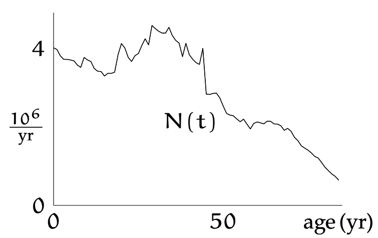

Comments on page 1

great file name

this seems like it should go with divide and conquer at the beginning of the semester

Read Sections 7.1 and 7.2 for Sunday's memo. The final (fifth) page has a fun problem to think about, which we'll look at on Monday. Have a nice weekend.

I am still not convinced or do I understand the whole concept of lumping?

I just got deja vu here...wasn't this line used in a previous reading?

I think so... and he's used it in class a few times.

Yah...he uses it a LOT in class. Sometimes its use is rather confusing, but it makes sense here.

Might as well call it, "As Sanjoy Mahajan says..."

Yeah this has been used a ton of times it's kind of a theme of the class

This seems like a very unspecific question and thus requires more difficulty in approximating...it would help to ask maybe how many babies are born in __ time, or how many people between 0-2 are in the US?

Um, I think you're reading too much into this. This is just the title, and a little subtitle to go alongside the main topic of estimating populations. The actual question will be addressed in the text.

That's the very first step that's taken in the paragraph: defining the parameters of the problem for estimation.

I've been asked this exact question as part of consulting case interviews before, I really like seeing examples like this because they are so applicable to what I want to do in the future!

True, but wasn't this method a little too simple and also very sensitive to the guess you made based on everyone dying at 75 years?

I'd say that the method is not at all too simple...esp. for a problem you encounter in interview where you need to thing quickly. Having a simple (but accurate) method is extremely valuable.

7
Lumping

7.1 Estimating populations: How many babies? 133
7.2 Bending of light 135
7.3 Quantum mechanics 138
7.4 Sound and electromagnetic radiation 143
7.5 Boundary layers 143

The symmetry chapter (Section 3.1) introduced the principle of invariance: 'When there is change, look for what does not change.' However, when you cannot find any useful but unchanging quantity, you have to make one. As Jean-Luc Picard often says, 'Make it so.'

7.1 Estimating populations: How many babies?

The first example is to estimate the number of babies in the United States. For definiteness, call a child a baby until he or she turns 2 years old. An exact calculation requires the birth dates of every person in the United States. This, or closely similar, information is collected once every decade by the US Census Bureau.

As an approximation to this voluminous data, the Census Bureau [33] publishes the number of people at each age. The data for 1991 is a set of points lying on a wiggly line $N(t)$, where t is age. Then

$$N_{\text{babies}} = \int_0^{2\,\text{yr}} N(t)\,dt. \qquad (7.1)$$

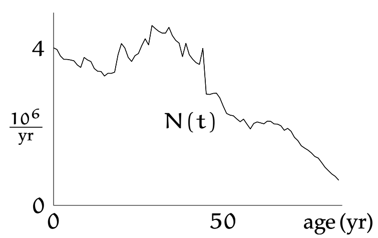

This is a good example problem.

the first example OF LUMPING is... you have yet to define what you are actually talking about in this chapter

This is a good point. I got to the end of the baby section and was thinking "So lumping is making integrals boxes?" I'd like to have a bit of an explanation of what is going to happen with lumping so I know what to look for in the coming sections.

I felt the same way. I read up to page 3 and realized I STILL didn't know what lumping was. I had an idea, but it was never specifically stated.

I think the introduction should define it immediately.

This just sounds very odd.

I would change "For definiteness" to "lets"

I agree this sounds strange.

this phrase is a little odd to me. perhaps try: "in order to do this, we must first define a range of ages for which the child is considered a baby"

Or you could just say, "For precision, ..."

I like both "for precision" and "lets". Alternately, we can word this as "Let's define a child as a baby..."

7
Lumping

7.1 Estimating populations: How many babies? 133
7.2 Bending of light 135
7.3 Quantum mechanics 138
7.4 Sound and electromagnetic radiation 143
7.5 Boundary layers 143

The symmetry chapter (Section 3.1) introduced the principle of invariance: 'When there is change, look for what does not change.' However, when you cannot find any useful but unchanging quantity, you have to make one. As Jean-Luc Picard often says, 'Make it so.'

7.1 Estimating populations: How many babies?

The first example is to estimate the number of babies in the United States. For definiteness, call a child a baby until he or she turns 2 years old. An exact calculation requires the birth dates of every person in the United States. This, or closely similar, information is collected once every decade by the US Census Bureau.

As an approximation to this voluminous data, the Census Bureau [33] publishes the number of people at each age. The data for 1991 is a set of points lying on a wiggly line $N(t)$, where t is age. Then

$$N_{\text{babies}} = \int_0^{2\,\text{yr}} N(t)\,dt. \quad (7.1)$$

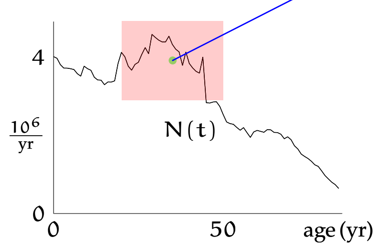

I like seeing these problems that are from all the way back when we did the diagnostic. It gives me a good sense of the progression of the class and makes me see how much I've learned when comparing to my old methods.

I agree. There were many problems on the diagnostic that had me absolutely stumped as to how to even start them.

But as we encounter tools that allow us to answer them, it's like a 'light bulb' moment.

I'm still looking forward to the golf ball problem!

Haha yeah, me too! I totally did this problem differently. I used the number of people in the US and dived and conquered to get to the number of babies (guessed how many were able to have children, and used the average of 2 children per couple).

Agree, I remember looking at the irregular spring Hooke's ;aw problem at the beginning of this class and being so confused. It was awesome being able to actually arrive at the answer this time around, you really feel like you're making progress.

Very relevant topic this time of year

Yep, I hope everyone has turned in their forms!

haha, this reminds me I need to fill in my after I finish my two tests this week.

This paragraph could use a better introduction, explaining what's going to be done with the data. I expected a simpler method to be explained first.

Is this the baby boom?

Well I guess the data is from 1991...

Quoted from next comment (posted @10:37):

For this data in particular, the sharp rise and steady-high numbers would represent the Baby Boomer generation. (the sharp jump just to the left of this box)

The sharp decline just shy of 50 would most likely be due to the number of lives lost during WWII and the Vietnam War.

Very interesting... You seem to be right

7

Lumping

7.1 Estimating populations: How many babies? 133
7.2 Bending of light 135
7.3 Quantum mechanics 138
7.4 Sound and electromagnetic radiation 143
7.5 Boundary layers 143

The symmetry chapter (Section 3.1) introduced the principle of invariance: 'When there is change, look for what does not change.' However, when you cannot find any useful but unchanging quantity, you have to make one. As Jean-Luc Picard often says, 'Make it so.'

7.1 Estimating populations: How many babies?

The first example is to estimate the number of babies in the United States. For definiteness, call a child a baby until he or she turns 2 years old. An exact calculation requires the birth dates of every person in the United States. This, or closely similar, information is collected once every decade by the US Census Bureau.

As an approximation to this voluminous data, the Census Bureau [33] publishes the number of people at each age. The data for 1991 is a set of points lying on a wiggly line $N(t)$, where t is age. Then

$$N_{\text{babies}} = \int_0^{2\,\text{yr}} N(t)\,dt. \qquad (7.1)$$

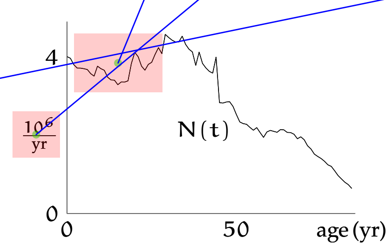

Why is there this dip and then rise in the number of people?

There are different numbers of people born every year, so it's possible for more people to be born in year x and then less people born in year x+k, which would account for the dip and then rise.

For this data in particular, the sharp rise and steady-high numbers would represent the Baby Boomer generation.

The sharp decline just shy of 50 would most likely be due to the number of lives lost during WWII and the Vietnam War.

It would be nice to see a date on when this data was taken - which census? That could explain the dim in people.

He does, it is the data for 1991 (so the 1990 Census)

true, but a title on the figure would also be nice.

When using data like a census? How far back is reasonable? 1991 was a decade ago, so using that data kind of seems unreasonable to me.

Guess he just couldn't find the figures from 2000. 1991 isn't so bad though.

These units kind of confuse me. I understand that if we take the integral we'll get the correct units, but does the census give the data as a rate. It seems more logical that they would just have the ages of people and a simple summation would do.

I agree, that is a bit weird. It makes the true nature of the graph look a little strange. Wouldn't the integral be different too if it is measured as a ratio?

Yeah I don't understand it either. The census takes data of how many people are currently in the country.

Is this necessary? We can see the chart on the side. Maybe reference that instead?

yeah i agree, this sounds a bit weird

I think addressing the graph itself would definitely be better. This is kind of out of context given that the rest of the text is professionally written and aimed at a fairly high level.

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

What are the height and width of this rectangle?

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \qquad (7.2)$$

Why did the life expectancy drop from 80 to 75 years?

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

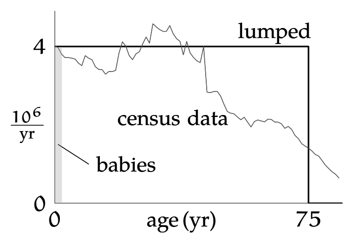

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \qquad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

Comments on page 2

When computing the area, obviously it works out when multiplying. Intuitively, we're really saying "number of people per age group" where age group is a year.

It's interesting, though, I don't think I would have noticed it unless this question pointed it out!

Does it have to do with looking at a particular year? It works out dimensionally when integrating (people/year * dt= people)

I think it's because the bin widths are small enough (1 day) that this is considered a smooth curve (and therefore why we integrate, instead of taking a discrete sum, above).

People per year is more useful than people per day, or just 'People' if this were presented as discrete data with each bin width equal to one day.

Also, "people" would be somewhat constant, as theoretically births equals deaths, and so we would get no useful information out of that graph

Mathematically, it's because we're integrating over the time period 0 to 2 years. We defined a baby to be a person between 0 and 2 years old, and since we're integrating over this time period to get a unit which is number of people, the vertical axis must be units of people per year.

Perhaps, having the vertical axis in units of people/year will specify the rate of change of people in a particular age group. This could be more useful for interpretation since it will allow us to see changes in the rate, which could help with anticipating future effects before actual changes happen.

You might want to specify "the vertical axis" in this sentence. For the longest time I was really confused as to why the horizontal axis supposedly had units of 1/T when it was labeled with units of time...

I got caught up in this same thing - based on the previous sentence it could even say "why does this axis..."

You might want to specify "the vertical axis" in this sentence. For the longest time I was really confused as to why the horizontal axis supposedly had units of 1/T when it was labeled with units of time...

That sounds redundant. If anything, just change "the axis" to "this axis"

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \quad (7.2)$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

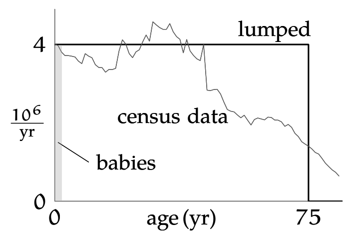

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \quad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

Yeah, I was actually thinking to myself while reading the above method that it was not practical, especially not in terms of approximations

Also, i'd specify what method is used. this is like the very precise standard method so having some sort of adjective, "this ___ method" would help

Yea I agree with that. When I originally approached the problem without reading, I thought that I could just do that and get the answer by subtracting they different y axis numbers each year. This is too simplistic.

You might want to explain what you mean by "desert island" and "back-of-the-envelope". I know what you mean coming from this class, but if other classes at other schools are going to be using this, then they might wonder why you're talking about desert islands.

I wonder why this would bother you if you were on a desert island to begin with. You'd think you'd rather spend time trying to get off the island than figuring out how many babies there are.

Got a lot of time on your hands I guess.

I think we've used these terms so many times throughout this class that mo matter where you come from if you've been paying attention they wouldn't bother you at this point.

This is a good explanation of the flaws of the brute force method.

This paragraph is awesome. It's one of those "this is why you're taking this course" paragraphs.

what does it mean to be integrated numerically

It just refers to finding the numerical solution of a particular curve

but we can simplify the integration using a similar, easy form

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \tag{7.2}$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

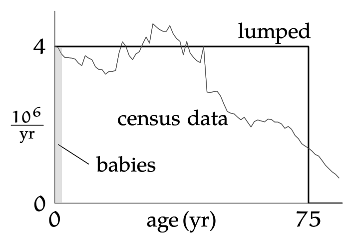

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \tag{7.3}$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

With today's computational tools, this isn't much of a problem–unless you're also referring to the "you're on a desert island" problem.

also, the idea of an exact numerical integration goes against the very idea of this class.

Numerical integration isn't really compatible with backs of envelopes.

I disagree. You can make a rough estimate by treating the trends in this graph as 3 lines...or even one plane and one slope

A rough approximation to the curve doesn't count as 'numerical integration', at least not in the sense he's using the term. http://en.wikipedia.org/wiki/Numerical_integ

well, wether or not it is much of a problem, it provides little insight and has minimal transfer value as stated, it's data specific, where it is about generality here.

I'm not sure how computing the number of babies can be translated to other problems, though...

agreed. I also think that integration in general, unless there is some sort of symmetry involved, isn't ideal when trying to approximate things.

why wouldn't you want to use data specific to this problem? I understand you're trying to teach something here, but this doesn't make sense to me.

I think this connects to the rest of his sentence about generality. When you are doing a integral, you want something that can be used for any problem not specifically a problem about babies. Also the integral may be significantly different in another time period. The 2010 census may yield data that will change the result and graph completely.

The focus of the class, and the methods we've been exploring, seems to be on portability and ease of reproducing methods of estimation. Data-specific solutions are the exact opposite of that and don't really teach you anything for your next encounter with a problem.

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \quad (7.2)$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

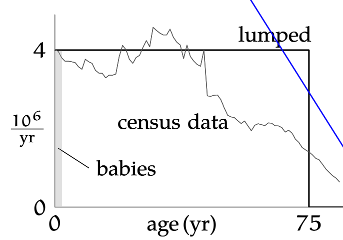

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \quad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

How can this have little insight when we compare our estimation this calculation for validity?

I believe this is mainly a reference to the point about generality.

General solutions (usually) provide more insight into the context surrounding the problem compared to exact solutions relevant only to specific data.

My problem with this statement is that integration is a very general technique for solving problems, and while this particular integral may not be useful for any other problems, neither will Sanjoy's lumping calculation, because both are particular to this problem. However both techniques can be widely applied to a number of problems with good results...

But see, that's not true. Integration is not a general technique, not in the slightest. When you are done with the problem, all you have learned is how to integrate that equation – which is great for solving the problem, but not useful for understanding the limits and framework of the quantities. The lumping technique can be applied to any integration, as well as various other uses I'm sure we'll cover as we explore the unit. They're just different classes of problem-solving.

Huh? If you're solving a problem with mathematics, shouldn't the math be specific to the problem?

I think he's saying that you want a way to solve the problem that doesn't matter on the specific data in the problem. You want to find a mathematical technique that will give you the solution of any data.

I agree with the second post, but I also think this sentence doesn't really help to illustrate his point. Integration is a very general procedure, and one that we could apply to this problem particularly by integrating the Census data.I think it would be more worthwhile to make a different case for why we might want to use a lumping technique instead of straight up numerical integration, since both are valid ways of approaching the problem.

I think what he's saying here is that in this case we're doing a "numerical integration" for this very specific curve. And we should be more interested in a general solution we can apply to more than just this exact situation. so he approximates the curve as one rectangle

the rectangle inherently has a large "transfer value" then?

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

What are the height and width of this rectangle?

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \tag{7.2}$$

Why did the life expectancy drop from 80 to 75 years?

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

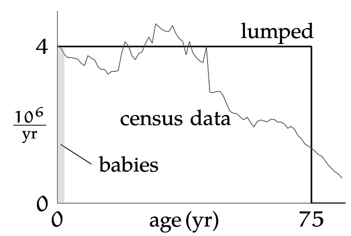

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \tag{7.3}$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

How does the estimation we did using lumping have transfer value? (Not that it doesn't, but I want to know what exactly is transferable to better understand why this was brought up as the third reason)

I think there is an easier way to go about the problem. It is common knowledge that the US population is roughly 300 million, the average household is about 3 people, distribute parents in each household by age distribution, and estimate how likely a couple in each age group will have a child that year...

answered my previous comment. I got a bit ahead

A rectangle? why specifically a rectangle? wouldn't it be more accurate as some sort of parabola curve?

Don't we lose the strong skew expressed by the data?

more accurate, yes. but a rectangle works well & is much easier to do the math for.

doesn't it kinda look like it should be a triangle instead- did you choose rectangle because it is so easy?

But this doesn't deal with the desert island problem either, although it helps with the integration.

Agreed, this solves the issue of not having integration tools on a "desert island", but you definitely wouldn't have the Census data and therefore any integral to try to approximate.

But it does help if we at least have a rough guess about the total US population and life expectancy (as he shows below). Then you can "integrate" your own estimated curve. I agree though that the wording in this section does make it sound like we're working directly off of the census data and not on our own estimates.

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \quad (7.2)$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

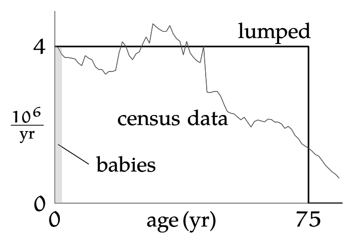

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \quad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

This is definitely an approximation method that we've all learned about from calculus

This approximation makes the assumption that there are equal numbers of people of each age though, correct?

Nope, it doesn't because its still using this ratio form...if I'm not mistaken.

I suppose if you're only looking at one section, you can choose the average of that section only as opposed to the average of the entire curve.

Right, the more sections that you lump, the more accurate your approximation will become. This is cool because this is exactly how we learned to do integrals in high school, taking the limit of the number of lumpings on a curves area.

Yea, I like how each of these methods we're learning all relate back to something that we learned in the beginning of our classes. I remember that the first thing we learned in Chemistry before the equations was solving using dimensional analysis and then this brings us back to calculus

It is a very simple, useful method.

I would assume you would make the initial value of the curve the height of the rectangle. right?

why are we using life expectancy rather than just looking at the region between 0-2 years?

We're using what we feel are two "known facts" to estimate something we don't know as well: the height of the rectangle.

The two knowns are: total population (area of the rectangle) and life expectancy (width of the rectangle)

Then we divide area by width to obtain height.

Now that we have height, we multiply it by 2 years to get the area spanned by babies, giving us # of babies.

To answer your question more, we ARE looking at the region between 0-2 years. But we are choosing to do so by first approximating the graph as a rectangle. To do that, we needed to do what I just said above.

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \tag{7.2}$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

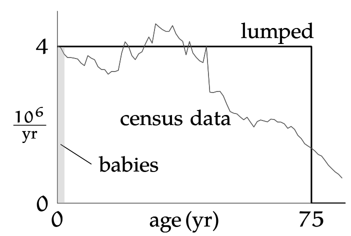

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \tag{7.3}$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

Since we use 75 anyway for the calculation later on, can't we just say that the life expectancy is roughly 75 years to avoid confusion?

Well, if he's trying to point out that you can fudge your numbers a bit, starting out with 75 doesn't provide a starting point for that.

thanks

so it seems like this works ok for the early ages, but if we were trying to estimate the number of elderly it might be way off.

Can we assume this? The curve slopes down pretty far by the time we reach 80

As mentioned above though that was due in large part to WWII, so today, the curve would not fall so abruptly at 50, it would fall at 60, and in the theoretical future, would fall at 80

I think we're basically figuring that 80 is right around the center of the curve (if we take life span as a sort of bell-like curve).

I would actually estimate it to be a little less.

I don't have a good sense of population growth in the US, so how long will we be able to use 300 million for future back-of-the-envelope calculations?

Should you consider how many babies a typical person has?

I think you could use that information as another method to get to a (hopefully) similar solution

Yes you could do it this way, but you'd have to think about what size the baby-having subsection of the population would be, how many babies each has on average, etc. These probably aren't as easy numbers to estimate as average life expectancy or US population.

I agree and the point of this exercise is to show us how to effectively lump things together, such as the graph.

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \qquad (7.2)$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

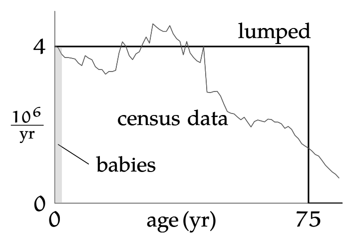

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \qquad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

is there a way other than this to estimate the height? I still don't fully understand what this unit represents.

The height could be seen as an estimate of the number of people at a given age. What's important is that by integrating the graph, we get a total population. Since the graph is represented as a rectangle, then a simplified assumption of height is area/x-axis.

It may be helpful to include this in the description of height and maybe even to describe why finding the height is important before we do it. It seems like a lot of people get confused when they aren't explicitly told why we're doing something .

I thought we were using 80 years?

I think he changed it when he realized that 80 doesn't divide 300 evenly.

It might be a good idea to check out case and point- it's a book about case interviewing that approaches this exact approximation in a slightly different manner, it might be useful to see an alternate- perhaps even less complicated way to do it.

It took me a while to understand what this was asking- I don't really think it needs to be addressed

I think it definitely does need to be addressed. In one sentence, he says the life expectancy is 80 years, but then he turns around and uses 75 years in the calculation. I think some explanation is necessary, otherwise I know I would be confused.

Yeah, this is very confusing, my understanding is it dropped to 75 just so that the math would work out nicer.

I think the wording of the question is a little awkward. Maybe "Why did we use 75 instead of 80" would work better?

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

What are the height and width of this rectangle?

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3\cdot 10^8}{75\,\text{yr}}. \qquad (7.2)$$

Why did the life expectancy drop from 80 to 75 years?

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4\cdot 10^6\,\text{yr}^{-1}$.

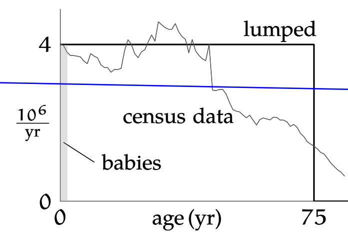

Integrating the population curve over the range $t = 0\ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4\cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8\cdot 10^6. \qquad (7.3)$$

The Census Bureau's figure is very close: $7.980\cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

I feel like this should be said earlier - "instead of 80 years, we will use 75 years because...".

Is this because we take out the 2 years for the babies? Or is it just to make our calculations easier.

I agree with the first person it isnt really clear what is going on here

I think its to make the resulting numbers easier to deal with. Also, It can't really be mentioned earlier, because until we know the population and see the division, we don't really know that 75 will produce a better number than 80.

Taking 2 years out for the babies would only make sense at all if that 300 million excluded children younger than 2 years old. This is just to make calculations easier.

We regards to the original comment: It may be difficult to move this explanation any earlier because he needs to set up the height calculation in order to explain why 75 years is more convenient.

I don't agree; it changed, and was immediately explained in the next line. I think that the small change is actually being used to teach another lesson here also (hence, the subsection for it).

we should have said this in the paragraph above, so that people dont get confused when they see the 75 as I was.

agreed. i didn't get confused but i didn't notice the change either. so it was more of a double-take

I agree as well - seems little reason to explain 80 if we're just going to change it around later.

or you could've used 81

It's a little hard to see that the babies rectangle is grayed in. Maybe make it darker?

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \tag{7.2}$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

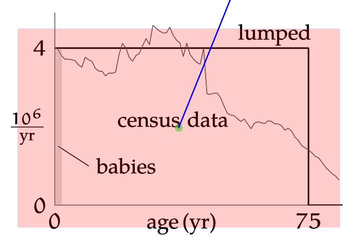

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \tag{7.3}$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

Note that you're plotting a calculation based on 2008 population on top of 1991 census data (when the population was 258 million – 14% less!).

Other than being bad practice, it makes it look like we chose an average that was too high for the whole population (it's pretty clear that the population doesn't integrate to be the area of the box).

Wow, good point. I was wondering why the box seemed to be so much larger than the curve it was supposed to approximate.

It's partly that, and could also be a couple of other things related to our life expectancy approximation.

First, there is some error in our estimate.

Second, using life expectancy as an estimator for population distribution ignores time in a couple of ways. Time has an influence in both changing population growth rates and changing life expectancies. Changing growth rates means more babies each year than the last.

An increasing life expectancy – and this is where it gets tricky – would mean that the true average life expectancy of the population (still measured from birth, but at the time of each person's birth, not in 1991 for everybody) would be lower than this estimate 75 or 80 years.

I don't, however, think these are significant factors, and I do think our analysis is a good approximation. Mostly I'm just bothered by the plotting of 2008 data on a 1991 graph, especially because (as in the Exxon example) a truer representation would help Sanjoy make his point even more convincingly.

You're right, the population in 1991 was less than in 2008. However, I think the reason he chooses to use the 2008 value of 300 million is because it's an approximation. If you actually used the 258 million value from 1991, you would probably round that up to 300 million to make the calculations easier anyway. Yes, it's 14% less, but I think in the approximation world, if it's in the same factor of 10, then it doesn't really matter.

A glance at the area covered by each shape makes it seem as though the rectangle should be a little smaller.

I had the same thought, it seems to work better with the same height but a width ending at age 70 or so instead.

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

What are the height and width of this rectangle?

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \quad (7.2)$$

Why did the life expectancy drop from 80 to 75 years?

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

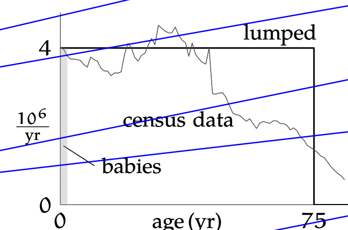

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \quad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

"and it might even cancel the lumping error": i feel like that stuff like this happens a lot in this class. and it's always a little too serendipitous for my comfort.

I somewhat agree...why we change the number makes sense, only I probably would have chosen 70 or 75 to start with, not 80. Numbers always magically working out make me think that the answer was already known...and now we just tweak things slightly so they fit

Although this kind of stuff happens a lot, there are also a fair number of examples where he says "the actual answer is 3 times greater" or "this approximation is off by a factor of 10 because.." which leads me to believe he's not using the answer for the approximating.

This is pretty true, but when we introduce error I'm still uncomfortable with the idea that "this error will hopefully cancel out the other error" without any explanation of how cutting 1/16 off of one value fixes the lumping error.

I agree that it's not ideal to try to make up for previous errors by making more error, but sometimes when we make a guess and we know we're low, but don't know by how much (as in we know we're low, but if we try to guess a higher value we might go over, etc), then it might make sense to try to overestimate something else.

Can you just make that generalization?

I feel like this sentence doesn't really add in terms of clarification.

I don't see how using 75 years "makes" the height 4*10^6 yr^-1.

Is their a logical way to get this on the y axis or is it just eyballing and guessing?

I think this problem is a very good example of "making it so".

I don't think making the box width 75 was necessary for this calculation

I agree. we should have chosen a box that fit the 0-2 year demographic better, and used that instead.

Yep, that seems unnecessary for htis demographic.

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \qquad (7.2)$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

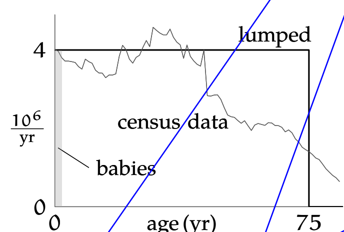

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \qquad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

I don't know if height is the right descriptor... (especially because infancy isn't labeled 'width', nor do I think it should be). I don't personally have a better suggestion, but this seemed off.

Maybe plotting the infancy 'slice' of the rectangle on the graph would make what this calculation is doing more clear.

Or maybe just calling it 'infancy width' or something.

that is really amazing how close it is, even looking at the curve there is a lot of error and it seems like the rectangular approximation would overestimate most ages by a lot, but it works well for babies. if we were trying to estimate the number of people between 50-60 would we use the same approximation or how would it change?

Why do these errors tend to cancel out so often?

It's important to understand what each error does to the problem. By lumping, we are counting the later years more then they should be but when we drop the age from 80 to 75, we are eliminating from the later years. It's important to know in which direction each error affects the answer.

This is impressive. I can think of a few other ways to estimate this number, but none could come nearly that close.

Does reducing the age range from 80 to 75 tend to add people just to the end of the distribution or is it more evenly distributed?

No matter how many times it happens, I am always amazed at the calculations ending up so close to the actual values, and I usually expect it each time at this point too.

is there a way to tell if the cancel each other out or multiply to make it worse?

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \tag{7.2}$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

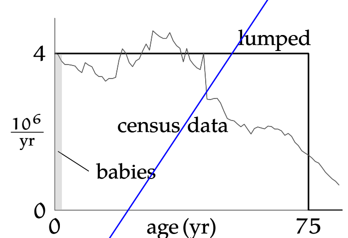

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \tag{7.3}$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

But couldn't the fudging have just as easily made the error greater? It feels like you just got lucky that the life expectancy you chose happened to cancel nicely with the lumping.

I think the decision to fudge the life expectancy downward was not random, but done for 2 reasons. One, to make division easier, and two, to cancel the lumping, which overestimated the population to 300 million.

There was some logic to picking a lower age and minimizing error. Looking at the distribution of the curve, it would be a mistake to use 85 instead, since you'd be making your approximation worse. Even without the plot, if you know that the average lifespan is 80, you want to pick something lower than that since the distribution is higher at the younger ages.

So basically lumping means making broad but reasonable generalizations?

This seems like one of the most basic forms of estimating, like back to the beginning. I don't quite see how its "lossy" though, since we aren't necessarily throwing things out, we just ignore them like we did at the outset of the class and with random constants always.

I saw it as "lossy" since we have ignored all the details presented by the U.S. Census. By showing the graph, we know that N(t) is some complex function. Thus, we ignore the complexities of the problem to form an estimation.

The errors canceling out almost perfectly seems like serendipity to me.

It's not entirely random; it's about finding the most important pieces of information and using them to find your solution in an easy way.

Yeah, I see lumping as using common sense approximations–what we've been doing at the beginning of the course. For example, estimating volume of a sphere as a cube for easy math. Except now, that sphere is some complex shape, and we have no choice but to approximate.

it is safe to assume that this rarely happens when applied to most problems though, correct?

Problem 7.1 Dimensions of the vertical axis
Why is the vertical axis labeled in units of people per year rather than in units of people? Equivalently, why does the axis have dimensions of T^{-1}?

This method has several problems. First, it depends on the huge resources of the US Census Bureau, so it is not usable on a desert island for back-of-the-envelope calculations. Second, it requires integrating a curve with no analytic form, so the integration must be done numerically. Third, the integral is of data specific to this problem, whereas mathematics should be about generality. An exact integration, in short, provides little insight and has minimal transfer value. Instead of integrating the population curve exactly, approximate it—lump the curve into one rectangle.

▶ *What are the height and width of this rectangle?*

The rectangle's width is a time, and a plausible time related to populations is the life expectancy. It is roughly 80 years, so make 80 years the width by pretending that everyone dies abruptly on his or her 80th birthday. The rectangle's height can be computed from the rectangle's area, which is the US population—conveniently 300 million in 2008. Therefore,

$$\text{height} = \frac{\text{area}}{\text{width}} \sim \frac{3 \cdot 10^8}{75\,\text{yr}}. \qquad (7.2)$$

▶ *Why did the life expectancy drop from 80 to 75 years?*

Fudging the life expectancy simplifies the mental division: 75 divides easily into 3 and 300. The inaccuracy is no larger than the error made by lumping, and it might even cancel the lumping error. Using 75 years as the width makes the height approximately $4 \cdot 10^6\,\text{yr}^{-1}$.

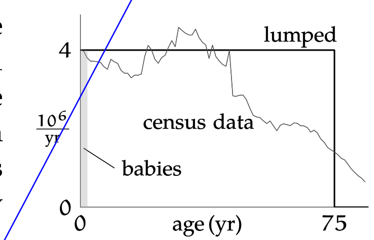

Integrating the population curve over the range $t = 0 \ldots 2\,\text{yr}$ becomes just multiplication:

$$N_{\text{babies}} \sim \underbrace{4 \cdot 10^6\,\text{yr}^{-1}}_{\text{height}} \times \underbrace{2\,\text{yr}}_{\text{infancy}} = 8 \cdot 10^6. \qquad (7.3)$$

The Census Bureau's figure is very close: $7.980 \cdot 10^6$. The error from lumping canceled the error from fudging the life expectancy to 75 years!

I like starting with this simple example... it makes the concept much easier to grasp without having to fight through a bunch of physics.

I agree, it's really given me a good idea of what lumping is all about too.

I agree. It was simple, clear, easy to follow, and helpful in understanding lumping.

Problem 7.2 Landfill volume
Estimate the US landfill volume used annually by disposable diapers.

Problem 7.3 Industry revenues
Estimate the annual revenue of the US diaper industry.

7.2 Bending of light

The fundamental principle of lumping is to replace a complex, changing process by a simpler, constant process. Let's apply the method far beyond mundane concerns about the number of babies, using lumping to revisit the bending of starlight by the sun. Using dimensional analysis and educated guessing (Section 5.4), we concluded that the bending angle is roughly GM/Rc^2, where R is the distance of closest approach (here, the radius of the sun), and M is the mass of the sun. Lumping provides a physical explanation for the same result; it thereby helps us make physical predictions (??).

So once again imagine a beam (or photon) of light that leaves a distant star. In its travels, it grazes the surface of the sun and reaches our eye. To estimate the deflection angle by using lumping, first identify the changing process. Here, the changing process is the angle that the light beam makes relative to its original, undeflected path; equivalently, the photon falls toward the sun as would a rock. This deflection angle increases slowly after the photon leaves the star, increasing most rapidly near the sun. Because the angle and position are changing, which means the downward gravitational force is changing, calculating the final deflection angle requires setting up and evaluating an integral – while carefully checking items in the integral such as the number of cosines and secants.

In contrast, the lumping approximation is much simpler. It pretends that the deflection is zero until the beam gets near the sun. Gravity, in this approximation, operates only near the sun. While the photon is near the sun, the approximation pretends further that the downward acceleration (toward the center of the sun) is a constant, rather than varying rapidly with position. Finally, once the beam is no longer near the sun, the deflection does not change.

Babies probably use a diaper a day, right? Or more?

Definitely more than 1/day.

Many more. I would guess about 4.

Being off by a factor of 4 in this case would be considered good by our standards, right?

These are really good problems, they build off the answer we just got so you have an easy place to start your estimation.

yeah i agree..also i've heard that the question about estimating the US diaper industry is often asked as an estimation problem during interviews!

After having solved the earlier problem, this problem seems relatively easy to solve- it should just involve mutiplication by a few factors.

This topic came back.

This should have been said much earlier

Agreed, I feel like in most sections, we dive right into some problem or example, and we don't spend any time describing the concept that we're trying to learn. This short explanation should have been put before the baby example.

Agreed.

Eh, I kind of like it here.. the baby estimating example is a nice way to flow into this so when he talks about lumping you have already been introduced to it

I also agree with the last opinion. It's nice to get a feel for the concept first with an example before defining it. Just like when you teach someone a card game: you jump right in and then teach the detailed rules as you go along.

I agree, this would have been a good sentence for the introduction - before we launched into any examples of lumping.

It's nice to see that this has been stated here though - I too think it should be said earlier, but it should not be removed. Repitition = Importance!

yeah i agree...I completely understood the first example but initially I didn't quite understand what lumping was exactly.

Problem 7.2 Landfill volume
Estimate the US landfill volume used annually by disposable diapers.

Problem 7.3 Industry revenues
Estimate the annual revenue of the US diaper industry.

7.2 Bending of light

The fundamental principle of lumping is to replace a complex, changing process by a simpler, constant process. Let's apply the method far beyond mundane concerns about the number of babies, using lumping to revisit the bending of starlight by the sun. Using dimensional analysis and educated guessing (Section 5.4), we concluded that the bending angle is roughly GM/Rc^2, where R is the distance of closest approach (here, the radius of the sun), and M is the mass of the sun. Lumping provides a physical explanation for the same result; it thereby helps us make physical predictions (??).

So once again imagine a beam (or photon) of light that leaves a distant star. In its travels, it grazes the surface of the sun and reaches our eye. To estimate the deflection angle by using lumping, first identify the changing process. Here, the changing process is the angle that the light beam makes relative to its original, undeflected path; equivalently, the photon falls toward the sun as would a rock. This deflection angle increases slowly after the photon leaves the star, increasing most rapidly near the sun. Because the angle and position are changing, which means the downward gravitational force is changing, calculating the final deflection angle requires setting up and evaluating an integral – while carefully checking items in the integral such as the number of cosines and secants.

In contrast, the lumping approximation is much simpler. It pretends that the deflection is zero until the beam gets near the sun. Gravity, in this approximation, operates only near the sun. While the photon is near the sun, the approximation pretends further that the downward acceleration (toward the center of the sun) is a constant, rather than varying rapidly with position. Finally, once the beam is no longer near the sun, the deflection does not change.

This is the intro to lumping, it shouldn't be in the second example.

It seems like this is the main principle for a lot of the things we are learning in this class.

This was shown earlier, but this really solidified it. It might be better to introduce it somewhere earlier

so from these examples seem to suggest that lumping is like using the riemann sum approach as opposed to the integral. They are both sums, although the integral approach has more accuracy, its easy to make bars and add them up as you do in a sum.

I like that analogy, that really helps me visualize what how lumping estimates things.

I like that you acknowledge that your target audience is nerds.

Mundane could just refer to the earthly nature of babies, not that they aren't important.

Either way, great wording in this paragraph.

I also like that this parallels and example of bending light from an earlier reading

will we be required to memorize this for the final? or will it be given

I highly, highly doubt he would ever require us to memorize something like that.

I am wondering the same thing. From what I've seen on the psets we'll have to memorize it. Unless we can use a cheat sheet.

It might be nice if you put the same diagram as you had in the earlier chapter repeated here since we read about that a while ago

I agree, it's always nice for these types of problems to get a picture. I think visualizing the problem in the correct manner is a huge part of being able to approximate.

well assuming this was in a book of some sort, you could always return to the section, since he did refer to it.

Problem 7.2 Landfill volume
Estimate the US landfill volume used annually by disposable diapers.

Problem 7.3 Industry revenues
Estimate the annual revenue of the US diaper industry.

7.2 Bending of light

The fundamental principle of lumping is to replace a complex, changing process by a simpler, constant process. Let's apply the method far beyond mundane concerns about the number of babies, using lumping to revisit the bending of starlight by the sun. Using dimensional analysis and educated guessing (Section 5.4), we concluded that the bending angle is roughly GM/Rc^2, where R is the distance of closest approach (here, the radius of the sun), and M is the mass of the sun. Lumping provides a physical explanation for the same result; it thereby helps us make physical predictions (??).

So once again imagine a beam (or photon) of light that leaves a distant star. In its travels, it grazes the surface of the sun and reaches our eye. To estimate the deflection angle by using lumping, first identify the changing process. Here, the changing process is the angle that the light beam makes relative to its original, undeflected path; equivalently, the photon falls toward the sun as would a rock. This deflection angle increases slowly after the photon leaves the star, increasing most rapidly near the sun. Because the angle and position are changing, which means the downward gravitational force is changing, calculating the final deflection angle requires setting up and evaluating an integral – while carefully checking items in the integral such as the number of cosines and secants.

In contrast, the lumping approximation is much simpler. It pretends that the deflection is zero until the beam gets near the sun. Gravity, in this approximation, operates only near the sun. While the photon is near the sun, the approximation pretends further that the downward acceleration (toward the center of the sun) is a constant, rather than varying rapidly with position. Finally, once the beam is no longer near the sun, the deflection does not change.

I'm not quite sure how it's a "physical" explanation. I don't really see the physical part...

I think it means physical like making the population into a box like the babies problem. Population curves don't really mean much at first glance in terms of the birth rate, but we know rectangles and have a good intuitive grasp for them.

Isn't lumping still a kind of an abstraction? Physical makes me think we are talking about something more literal.

Definitely have no idea how this is possible.

are you saying that we are deriving the same result again or using it in other ways (i.e. lumping) to show that that is correct?

What are these doing here?

Probably a missing reference (citation, or possibly figure). This is how my missing references show up in latex documents.

I don't really understand what "changing process" means here.

I see what you're saying here, but I don't think the sentence is very clear.

Agreed, perhaps a diagram would show this better. I still have trouble thinking about the starlight grazing the sun with and without the pull of the sun. I guess it's because if I imagine light that just grazes the sun on its way to me, but consider the pull of the sun, then that light no longer just grazes the sun, instead (I think) it gets pulled in more so I wouldn't see it.

I see what you're saying here, but I don't think the sentence is very clear.

You usually break up your blocks of text with diagrams or figures. Your writing is easy to follow, but when I read on a computer screen large chunks of text are hard to process.

agreed. i'd like a figure somewhere on the page. you have two on the next page, maybe one can be brought up?

I feel like a diagram here would be really helpful. I'm a little lost as to what variables we're dealing with here.

Problem 7.2 Landfill volume
Estimate the US landfill volume used annually by disposable diapers.

Problem 7.3 Industry revenues
Estimate the annual revenue of the US diaper industry.

7.2 Bending of light

The fundamental principle of lumping is to replace a complex, changing process by a simpler, constant process. Let's apply the method far beyond mundane concerns about the number of babies, using lumping to revisit the bending of starlight by the sun. Using dimensional analysis and educated guessing (Section 5.4), we concluded that the bending angle is roughly GM/Rc^2, where R is the distance of closest approach (here, the radius of the sun), and M is the mass of the sun. Lumping provides a physical explanation for the same result; it thereby helps us make physical predictions (??).

So once again imagine a beam (or photon) of light that leaves a distant star. In its travels, it grazes the surface of the sun and reaches our eye. To estimate the deflection angle by using lumping, first identify the changing process. Here, the changing process is the angle that the light beam makes relative to its original, undeflected path; equivalently, the photon falls toward the sun as would a rock. This deflection angle increases slowly after the photon leaves the star, increasing most rapidly near the sun. Because the angle and position are changing, which means the downward gravitational force is changing, calculating the final deflection angle requires setting up and evaluating an integral – while carefully checking items in the integral such as the number of cosines and secants.

In contrast, the lumping approximation is much simpler. It pretends that the deflection is zero until the beam gets near the sun. Gravity, in this approximation, operates only near the sun. While the photon is near the sun, the approximation pretends further that the downward acceleration (toward the center of the sun) is a constant, rather than varying rapidly with position. Finally, once the beam is no longer near the sun, the deflection does not change.

i don't understand this phrasing.

Me either.

I think he means "also assumes"

This seems to me pretty intuitive, its the way we learn everything, first lumped into the basics, then gradually with more and more variables

Although this makes sense to me, it still hasn't reasoned far enough for me to know how to complete the problem.

is this different than the assumptions from last time- I am just remembering the drawing on the board or did it not matter since we were looking at dimensions?

From the diagram it makes it seem as though the force of gravity acts once, to deflect the point, and then stops acting on it. Am I misunderstanding that?

Nevermind, it's explained later.

Isn't this (using a constant) how we use gravity on earth?

I like this paragraph. It makes the similarity between how we're solving this problem and how we solved the last problem very clear - we're simply making a rectangle.

I agree...the contrast helps a lot in these building examples comparing multiple methods.

agreed, as an engineer, I really appreciated this paragraph.

I'm having trouble seeing how lumping is different from just making general approximations

It feels like the same approximations we did when we divided and conquered things like the MIT budget–there we used approximations that were lossy in the sense that we neglected or pooled things.

They definitely seem very similar.

yeah I agree..would it be possible to include an example in the text where the difference between divide&conquer and lumping is more noticeable? or are they always this similar?

Problem 7.2 Landfill volume
Estimate the US landfill volume used annually by disposable diapers.

Problem 7.3 Industry revenues
Estimate the annual revenue of the US diaper industry.

7.2 Bending of light

The fundamental principle of lumping is to replace a complex, changing process by a simpler, constant process. Let's apply the method far beyond mundane concerns about the number of babies, using lumping to revisit the bending of starlight by the sun. Using dimensional analysis and educated guessing (Section 5.4), we concluded that the bending angle is roughly GM/Rc^2, where R is the distance of closest approach (here, the radius of the sun), and M is the mass of the sun. Lumping provides a physical explanation for the same result; it thereby helps us make physical predictions (??).

So once again imagine a beam (or photon) of light that leaves a distant star. In its travels, it grazes the surface of the sun and reaches our eye. To estimate the deflection angle by using lumping, first identify the changing process. Here, the changing process is the angle that the light beam makes relative to its original, undeflected path; equivalently, the photon falls toward the sun as would a rock. This deflection angle increases slowly after the photon leaves the star, increasing most rapidly near the sun. Because the angle and position are changing, which means the downward gravitational force is changing, calculating the final deflection angle requires setting up and evaluating an integral – while carefully checking items in the integral such as the number of cosines and secants.

In contrast, the lumping approximation is much simpler. It pretends that the deflection is zero until the beam gets near the sun. Gravity, in this approximation, operates only near the sun. While the photon is near the sun, the approximation pretends further that the downward acceleration (toward the center of the sun) is a constant, rather than varying rapidly with position. Finally, once the beam is no longer near the sun, the deflection does not change.

would small angle approximations be considered a type of lumping? i feel like it would be, and i wanted to see if i'm understanding "lumping" correctly.

also. i think using biot numbers to determine whether or not the lumped parameter model can be applied, from 2.005, would be lumping as well?

Yeah. It's definite used in this example later. And it definitely feels like it belongs in this unit.

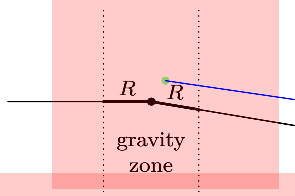

The problem then becomes one of estimating the deflection produced by gravity while the beam is in the gravity zone. But what is the near zone? As a reasonable guess, define 'near' to mean, 'Within R on either side of the location of closest approach.' The justification is that the distance of closest approach, which is R, is the only length in the problem, so the size of the near zone must be a dimensionless constant times R.

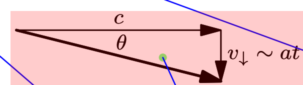

The deflection calculation is easiest at the location of closest approach, so assume that the bending happens there and only there – in other words, the beam's track has a kink rather than changing its direction smoothly. At the kink, the gravitational acceleration, which is all downward, is $a \sim GM/R^2$. The downward velocity is the acceleration multiplied by the time in the gravity zone. The zone has length $\sim R$, so the time is $t \sim R/c$. Thus the downward velocity is $\sim GM/Rc$. The deflection angle, in the small-angle approximation, is the downward velocity divided by the forward velocity. Therefore,

$$\theta \sim \frac{GM/Rc}{c} = \frac{GM}{Rc^2}. \tag{7.4}$$

The lumping argument has reproduced the result of dimensional analysis and guessing.

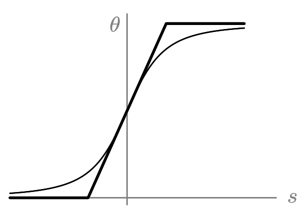

The true curve of θ versus position (measured as distance from the point of closest approach) varies smoothly but, as mentioned, it is difficult to calculate. Lumping replaces that smooth curve with a piecewise-straight curve that reflects the behaviors in and out of the gravity zone: no change in θ outside the gravity zone, and a constant rate of change in θ inside the gravity zone (with the rate set by the rate at the closest approach). Lumping is a complementary method to dimensional analysis. Dimensional analysis is a mathematical argument, although the guessing added a bit of physical reasoning. Lumping removes as much mathematical complexity as possible, in order to focus on the physical reasoning. Both approaches are useful!

...the crooked shall be made straight, and the rough places plain. (Isaiah 40:4)

this graph makes it much more clear

Could you include the sun in this picture with all of the places that R should appear? I think that would make the paragraph easier to understand. I had to read it a couple of times to decide what you meant by "R."

I agree. I thought the 'dot' was supposed to be the sun at first, as if the light was bouncing off the sun like a billiard ball. But then i realized it would have to be horribly out of scale...

I agree as well. I think a full diagram of sun, R, gravity zone and where you are measuring the deflection from.

Is it always best to use numbers we have for numbers that are arbitrary. I like it, keeps things simple.

This approximation still seems really arbitrary to me, although I guess it depends on how much it affects the final solution

I agree. This seems like a random choice.

I don't think it's that random, R is a property of the planet or star we are looking at so clearly it will have an effect on our gravity zone. Also by using R we can easily do some scaling later if our results do not seem to make sense

GM/cˆ2 is also a length. I would argue that this is also important. The slower something is moving, for example, the further away 'significant' deflections would start to accumulate.

True...except I believe we are talking about photons, so they are all moving at the speed of light

So? GM/cˆ2 being invariant in this problem doesn't make it irrelevant. Granted it is much smaller than R (since it equals R*theta and theta is small), but that's not particularly convincing, I don't think.

I like this diagram. It helps me organize my thoughts quite a bit.

Yeah, it's interesting using a triangle where the units are all velocities instead of distances.

Yea, you can use triangles for many different units. Aerospace uses velocity triangles a lot for trajectories and anything w/ relative velocities and positions when compared to things such as the wind or other disturbances

The problem then becomes one of estimating the deflection produced by gravity while the beam is in the gravity zone. But what is the near zone? As a reasonable guess, define 'near' to mean, 'Within R on either side of the location of closest approach.' The justification is that the distance of closest approach, which is R, is the only length in the problem, so the size of the near zone must be a dimensionless constant times R.

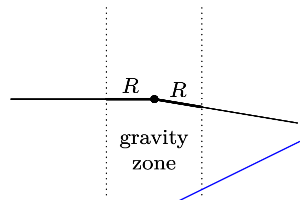

The deflection calculation is easiest at the location of closest approach, so assume that the bending happens there and only there – in other words, the beam's track has a kink rather than changing its direction smoothly. At the kink, the gravitational acceleration, which is all downward, is $a \sim GM/R^2$. The downward velocity is the acceleration multiplied by the time in the gravity zone. The zone has length $\sim R$, so the time is $t \sim R/c$. Thus the downward velocity is $\sim GM/Rc$. The deflection angle, in the small-angle approximation, is the downward velocity divided by the forward velocity. Therefore,

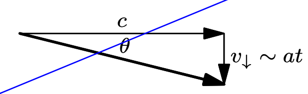

$$\theta \sim \frac{GM/Rc}{c} = \frac{GM}{Rc^2}. \quad (7.4)$$

The lumping argument has reproduced the result of dimensional analysis and guessing.

The true curve of θ versus position (measured as distance from the point of closest approach) varies smoothly but, as mentioned, it is difficult to calculate. Lumping replaces that smooth curve with a piecewise-straight curve that reflects the behaviors in and out of the gravity zone: no change in θ outside the gravity zone, and a constant rate of change in θ inside the gravity zone (with the rate set by the rate at the closest approach). Lumping is a complementary method to dimensional analysis. Dimensional analysis is a mathematical argument, although the guessing added a bit of physical reasoning. Lumping removes as much mathematical complexity as possible, in order to focus on the physical reasoning. Both approaches are useful!

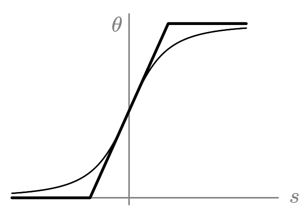

. . . the crooked shall be made straight, and the rough places plain. (Isaiah 40:4)

So i guess the "lumping" comes from lumping the total angle of deflection that is experienced over time into one angle that happens instantaneously. This makes it trigonometrically easier to evaluate.

Yeah I think that's correct.

Lumping the curve into a triangle

it would I think make more sense to put the diagram down here, so people don't wonder why there is a single kink instead of a smooth curve

isn't the diagram the one above? it clearly shows a bent line.

I don't think that is the diagram above. I am also a bit confused by the wording in general.

Seems like a lot of this is just proportional reasoning. Looking forward to seeing more problems that really rely on lumping.

this is a little confusing

Why don't you say length 2R when you know that's closer the actual length the light travels through the zone (it can't be anything less)? I understand we're concerned with dropping unnecessary complexity, but a factor of two doesn't seem too complex to me, and it might get us closer to the actual answer?

Although I'm still not quite sure on the whole lumping thing, this is a good explanation of 7.2, and a nice connection to previous things we've done. Very easy to understand.

not totally following this logic

The problem then becomes one of estimating the deflection produced by gravity while the beam is in the gravity zone. But what is the near zone? As a reasonable guess, define 'near' to mean, 'Within R on either side of the location of closest approach.' The justification is that the distance of closest approach, which is R, is the only length in the problem, so the size of the near zone must be a dimensionless constant times R.

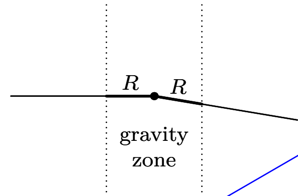

The deflection calculation is easiest at the location of closest approach, so assume that the bending happens there and only there – in other words, the beam's track has a kink rather than changing its direction smoothly. At the kink, the gravitational acceleration, which is all downward, is $a \sim GM/R^2$. The downward velocity is the acceleration multiplied by the time in the gravity zone. The zone has length $\sim R$, so the time is $t \sim R/c$. Thus the downward velocity is $\sim GM/Rc$. The deflection angle, in the small-angle approximation, is the downward velocity divided by the forward velocity. Therefore,

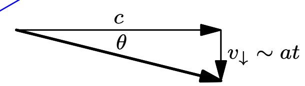

$$\theta \sim \frac{GM/Rc}{c} = \frac{GM}{Rc^2}. \tag{7.4}$$

The lumping argument has reproduced the result of dimensional analysis and guessing.

The true curve of θ versus position (measured as distance from the point of closest approach) varies smoothly but, as mentioned, it is difficult to calculate. Lumping replaces that smooth curve with a piecewise-straight curve that reflects the behaviors in and out of the gravity zone: no change in θ outside the gravity zone, and a constant rate of change in θ inside the gravity zone (with the rate set by the rate at the closest approach). Lumping is a complementary method to dimensional analysis. Dimensional analysis is a mathematical argument, although the guessing added a bit of physical reasoning. Lumping removes as much mathematical complexity as possible, in order to focus on the physical reasoning. Both approaches are useful!

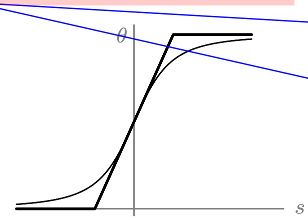

...the crooked shall be made straight, and the rough places plain. (Isaiah 40:4)

Where did this equation come from? For me it's a lot easier to follow these descriptions of mathematical processes if you list them as a series of equations like you do with 7.4 below and reference the equations that you used to derive the one you're listing.

this just comes from velocity = distance / time (and so time =distance /time). The distance is R, and c is the speed of light, so time = R/c.

Yea, this is just a basic velocity equation which I don't think needs a series of equations to lead up to it, but it may help to reiterate that c is speed of light just so that people know this is a velocity.

This t=R/c comes from the usual velocity equation. we have to find this because we know acceleration a=GM/(R^2) and velocity is acceleration*time. so by multiplying the 2, we can get velocity

You might want to mention that the small-angle approximation you're using is tan(O) (O). Most people will think of cosine when hearing small angle approximation.

small angle approximation: sin(x) x, cos(x) 1. It follows that tan(x) x.

it might not necessarily be intuitive that the deflection angle will be small

I agree, I wouldn't have jumped to the tan approximation by reading that.

Still good to point out. NB proves its usefulness here.

I had to re-read this section a couple of times, but it was well worth it!

This section was very clear and I always love to see a previous section revisited with a different method of approximation.

This is actually pretty cool.

I think it's important to note that these guesses are random and so we won't always get results that are the same as in other methods.

That's a good point...seeing as we have the answer already, lumping made this problem seem way simple. However, I feel like lumping would prone me to more errors

It is cool, and a good note of the robustness of these methods. Redundancy at its best. ;)

I like that with different methods, you can still reach the same conclusion in this class.

The problem then becomes one of estimating the deflection produced by gravity while the beam is in the gravity zone. But what is the near zone? As a reasonable guess, define 'near' to mean, 'Within R on either side of the location of closest approach.' The justification is that the distance of closest approach, which is R, is the only length in the problem, so the size of the near zone must be a dimensionless constant times R.

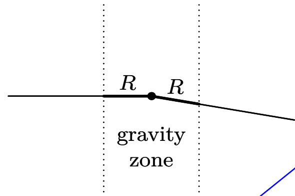

The deflection calculation is easiest at the location of closest approach, so assume that the bending happens there and only there – in other words, the beam's track has a kink rather than changing its direction smoothly. At the kink, the gravitational acceleration, which is all downward, is $a \sim GM/R^2$. The downward velocity is the acceleration multiplied by the time in the gravity zone. The zone has length $\sim R$, so the time is $t \sim R/c$. Thus the downward velocity is $\sim GM/Rc$. The deflection angle, in the small-angle approximation, is the downward velocity divided by the forward velocity. Therefore,

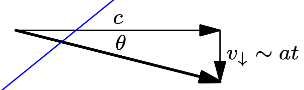

$$\theta \sim \frac{GM/Rc}{c} = \frac{GM}{Rc^2}. \tag{7.4}$$

The lumping argument has reproduced the result of dimensional analysis and guessing.

The true curve of θ versus position (measured as distance from the point of closest approach) varies smoothly but, as mentioned, it is difficult to calculate. Lumping replaces that smooth curve with a piecewise-straight curve that reflects the behaviors in and out of the gravity zone: no change in θ outside the gravity zone, and a constant rate of change in θ inside the gravity zone (with the rate set by the rate at the closest approach). Lumping is a complementary method to dimensional analysis. Dimensional analysis is a mathematical argument, although the guessing added a bit of physical reasoning. Lumping removes as much mathematical complexity as possible, in order to focus on the physical reasoning. Both approaches are useful!

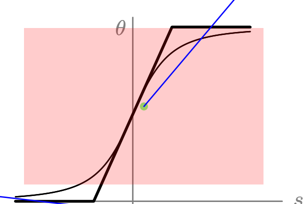

. . . the crooked shall be made straight, and the rough places plain. (Isaiah 40:4)

So accuracy in this case depends on what we choose for R, correct?

For the way we modeled the problem by neglecting the angle not near the sun. However how accurate can we get with a "better" R?

I'm not really sure, for me it's hard to conceptualize how to determine the accuracy in a lot of these approximations.

It also depends on the 'zone' we picked, for where gravity matters.

It was initially very unclear to me what this section meant – it took several readings to understand that this was position *along* the trajectory of the light, rather than some other measure of position.

I had to go back and reread a couple sections in this reading to appreciate what they were saying.

Looking at this diagram - this approach seems super useful for anything symmetric where the over and under estimates will cancel nicely.

This is a useful diagram. I think the addition of dashed lines at the angles would be nice. (dashed vertical lines).

I agree, this graph definitely helps visualize the process of lumping and how well it works in approximations.

Absolutely. The darkened line is a good attempt, but some dashes would be even better.

I agree–I think this diagram really shows how lumping can replace something complex (the smooth curve) with something simple (the straight curve)–this especially helps in problems like this where there are three distinct zones (outside the gravity zone, then inside, then out again), since each of these zones is in a sense separated. reminds me of divide and conquer!

This method of using piecewise straight curves to approximate quadratic looking curves is actually used a lot in my EE classes (bode plots, etc.) I guess those methods were kind of using lumping all along.

The problem then becomes one of estimating the deflection produced by gravity while the beam is in the gravity zone. But what is the near zone? As a reasonable guess, define 'near' to mean, 'Within R on either side of the location of closest approach.' The justification is that the distance of closest approach, which is R, is the only length in the problem, so the size of the near zone must be a dimensionless constant times R.

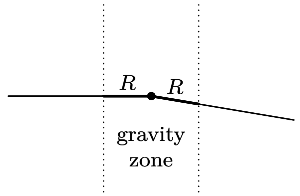

The deflection calculation is easiest at the location of closest approach, so assume that the bending happens there and only there – in other words, the beam's track has a kink rather than changing its direction smoothly. At the kink, the gravitational acceleration, which is all downward, is $a \sim GM/R^2$. The downward velocity is the acceleration multiplied by the time in the gravity zone. The zone has length $\sim R$, so the time is $t \sim R/c$. Thus the downward velocity is $\sim GM/Rc$. The deflection angle, in the small-angle approximation, is the downward velocity divided by the forward velocity. Therefore,

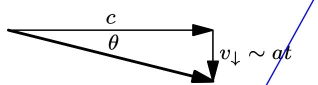

$$\theta \sim \frac{GM/Rc}{c} = \frac{GM}{Rc^2}. \tag{7.4}$$

The lumping argument has reproduced the result of dimensional analysis and guessing.

The true curve of θ versus position (measured as distance from the point of closest approach) varies smoothly but, as mentioned, it is difficult to calculate. Lumping replaces that smooth curve with a piecewise-straight curve that reflects the behaviors in and out of the gravity zone: no change in θ outside the gravity zone, and a constant rate of change in θ inside the gravity zone (with the rate set by the rate at the closest approach). Lumping is a complementary method to dimensional analysis. Dimensional analysis is a mathematical argument, although the guessing added a bit of physical reasoning. Lumping removes as much mathematical complexity as possible, in order to focus on the physical reasoning. Both approaches are useful!

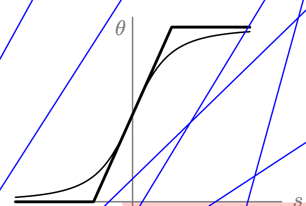

. . . the crooked shall be made straight, and the rough places plain. (Isaiah 40:4)

is gravity zone same as the near zone, beyond R?

Yes, it is the zone within distance R, where we consider gravity to have a meaningful effect on the photon.

lumping seems to just ignore units of measure instead of combining them in useful ways

Since you describe two regions of the graph, why not say the third as well? "At the end, theta remains constant, so the photon is traveling in a straight line but at its deflected angle."

I would like to see the results of this calculation involving lumping compared to the actual results of being integrated.

It almost feels like cheating. Almost.

Ha, yeah, it does. I wish I tried this more often during 8.01...

I've started using it in some of my other classes to get a feel for the answer before I really try it, and its really helpful!

weird switch in tense here

Yeah, grammatical correctness says it should be "adds."

I like this explanation.

Good summation of the tools and there realtions

yeah, these help me to cement the concepts that I've read and see how the over-arching theme connects with particular cases and examples. I think these summations are really useful and retaining the information.

I agree, I think you should set aside sentences like this that give a good summation of the methods.

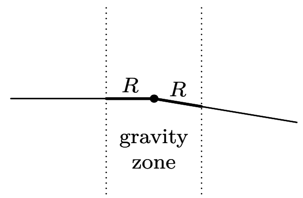

The problem then becomes one of estimating the deflection produced by gravity while the beam is in the gravity zone. But what is the near zone? As a reasonable guess, define 'near' to mean, 'Within R on either side of the location of closest approach.' The justification is that the distance of closest approach, which is R, is the only length in the problem, so the size of the near zone must be a dimensionless constant times R.

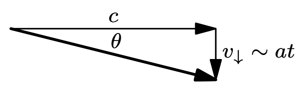

The deflection calculation is easiest at the location of closest approach, so assume that the bending happens there and only there – in other words, the beam's track has a kink rather than changing its direction smoothly. At the kink, the gravitational acceleration, which is all downward, is $a \sim GM/R^2$. The downward velocity is the acceleration multiplied by the time in the gravity zone. The zone has length $\sim R$, so the time is $t \sim R/c$. Thus the downward velocity is $\sim GM/Rc$. The deflection angle, in the small-angle approximation, is the downward velocity divided by the forward velocity. Therefore,

$$\theta \sim \frac{GM/Rc}{c} = \frac{GM}{Rc^2}. \tag{7.4}$$

The lumping argument has reproduced the result of dimensional analysis and guessing.

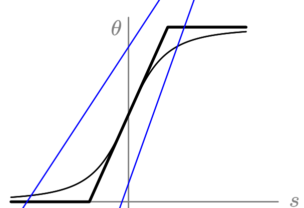

The true curve of θ versus position (measured as distance from the point of closest approach) varies smoothly but, as mentioned, it is difficult to calculate. Lumping replaces that smooth curve with a piecewise-straight curve that reflects the behaviors in and out of the gravity zone: no change in θ outside the gravity zone, and a constant rate of change in θ inside the gravity zone (with the rate set by the rate at the closest approach). Lumping is a complementary method to dimensional analysis. Dimensional analysis is a mathematical argument, although the guessing added a bit of physical reasoning. Lumping removes as much mathematical complexity as possible, in order to focus on the physical reasoning. Both approaches are useful!

...the crooked shall be made straight, and the rough places plain. (Isaiah 40:4)

Tell that to the math department when they take off points for "non-rigorous" thinking and reasoning...

that's why it's the math department! :)

this class is the art of approximation *in science and engineering*, where such approaches are not only allowed, but necessary!

Agreed! I think the math department wants to show the rigor behind such methods, so using the methods to prove the methods is a little silly.

I like this quotation, it's simple and humorous, yet speaks volumes about exactly how this class has taught us to approach problems.

I don't really like having scripture quotes in my textbooks, however.

I'd go further and say that I _really_ don't like having scripture quoted in non-fiction books not about religion

I love scripture quoted in non-fiction books about religion, I find it really funny and amusing; but then again I am not Christian.

This section reminded me of an interesting photon bending example that we see on a daily basis, although different principles are in effect. Every time the sun sets or rises, there is not yet a linear path from sun to that point on earth. You could estimate how much time there is between when light first reaches the earth and when there actually is a linear path.

Problem 7.4 Higher values of GM/Rc^2
When GM/Rc^2 is no longer small, strange things happen. Use lumping to predict what happens to light when $GM/Rc^2 \sim 1$.

Comments on page 5

It seems like this might result in the photon orbiting (or nearly orbiting) the incredibly massive object, rather than simply being deflected.

So you get a big swirl of light going around the sun?

I hope we go over this in class!

It is one of the definite items on the agenda for today.

I'll try this. Are we going to do this example in class? I think it would be interesting to go through it.

G=6e-11, C=3e8,Cˆ2=1e17, Schwarzschild radius of earth is about 1e-2, mass of earth=6e24

GM/Rcˆ2=6e-28(M/R)=1

if theta gets large it forms a black hole. Was this the answer they are looking for?

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

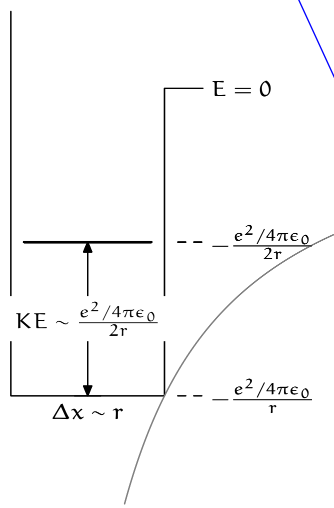

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

Comments on page 1

I JUST WANTED TO SAY:

I USED A DIMENSIONAL ANALYSIS on a complicated physics problem on waves today ON MY EXAM! I forgot the exact equation and was running low on time, didn't want to derive it, so I quickly pulled out dimensional analysis and did the problem in under 30 seconds. Thanks Sanjoy!

Just want to reiterate again how super sweet that was today on the exam!

-OP

if nb had a "like" button, i'd use it here :)

Read Section 7.3 for Tuesday's memo. It revists the hydrogen calculation from Reading 19 (r19-dimensions-hydrogen.pdf on NB), using lumping and easy cases to give a physical explanation of the earlier dimensional-analysis result.

Maybe you could reference the section in which we calculated this for the first time? It would be nice for the readers to be able to easily flip back to see the alternate method.

Agreed, but I also really like the fact that you're repeating earlier examples, just to show that you get similar answers, and as a check on your earlier calculations

Perhaps a brief statement would be nice, but it's best to keep it short. Again, this is a textbook so people can just look back. Maybe just a statement of which section it was in and also the approach used there, in contrast to lumping.

A lot of times he just uses an in-text citation like (1.2.3), I think that would be extremely helpful here.

I agree. I think a brief reminder would be helpful. I couldn't remember what t his was referring to immediately.

I remember the section vaguely, but mostly remember that I didn't understand it fully.

I tried to do that in the introductory note on NB (which says that reading 19 is the last example of it).

is this specific to different elements? or is it the same for all. and if it is specific does it just measure to the first electron shell or the outermost in the lowest energy state

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

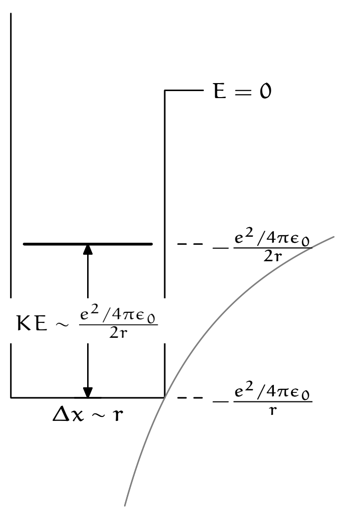

I feel like this part covers a lot of the things we've learned in this class.

how does d&c relate energy to the bohr radius distance?

D&C is used to analyze the two types of energy. The energy is already related to the Bohr radius as it is the radius of the orbit with lowest energy.

I'm assuming he will connect the total energy to the energy of the ground state. I also think that he was talking about using divide and conquer to figure out what the total energy (potential plus kinetics) was.

I remember an approach similar to this that was explored in my freshman seminar. It seemed to work out well so this seems like a plausible explanation

Is this a divide and conquer approach to the problem or are we simply dividing and conquering by calculating the potential and kinetic energy separating? I recall in divid and conquer we didn't always calculate all the variables involved but sometimes we just approximated with the ones that were greater in value.

This question confuses me – is it asking how one can understand it or where the physical entity that is kinetic energy originates in the universe?

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

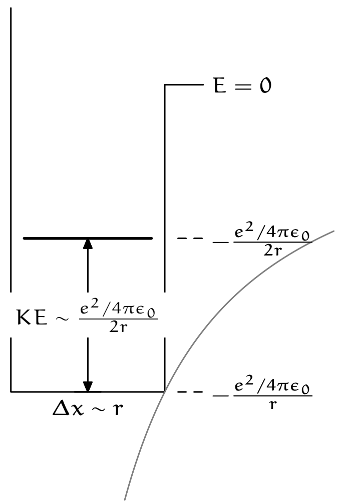

I don't know much about this stuff – is it in fact accelerating? Why?

is the reading implying that electrons do not accelerate? if so, why is this true?

That's the whole reason that quantum mechanics is necessary. Note the following series of classically true statements:

A body orbiting another body is constantly accelerating (since its velocity is always changing direction). Any accelerating charge emits electromagnetic radiation. Electromagnetic radiation is a form of energy. Energy is conserved. Thus as an electron orbits an atom, it emits (and thus loses) energy. So the electron will eventually run out of energy, spiraling into the nucleus and destroying the atom. But it doesn't (as evidenced by the fact that we're all still here to talk about it). QM provides a way to explain how an electron can be localized in an atom and yet not be constantly emitting energy.

that was a very good explanation. Not overtly tedious but detailed enough to get a general view of things. Thanks

great explanation

Also, can you explain why this would happen? Not sure I remember enough to rationalize it myself...

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

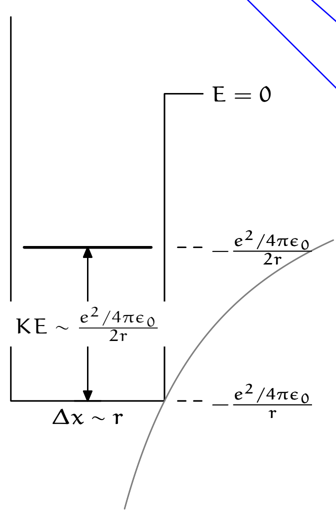

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

The fact that it didn't bothered me when I first encountered the idea of an atom in middle school.

I've never heard this explanation before. But after hearing it this whole concept makes much more sense.

In middle school? Wow.

Yeah, that's rather impressive. Especially considering this fact has never bothered me until now (well, technically a few readings ago when this issue was first brought up).

wow you must have been an incredibly smart and knowledged middle schooler to be on the verge of discovering quantum mechanics!

Not really. It was featured on Bill Nye the Science Guy, and my 8th grade science teacher gave an age appropriate introduction to how electrons move around the nucleus of atoms. No rediscovery of quantum mechanics needed.

oh. nevermind then.

Bill Nye is the man... just saying

Yeah, my sister liked Beakman's world, but I always enjoyed Bill Nye better; Beakman was a bit too hipsterish and something was just not right about him. Bill's always been a real bro, the man's man, something I can relate to.

Older sister? I, personally, learned a lot more from beakman...better science.

sarcasm much? but this is an interesting thing to point out; you would think they would mention it when going over the "electron cloud"

This seems depressing. You ask a question, realize that you don't know the answer, or that anyone else does either.

Pretty cool. I'd never heard it put that way.

'r' is only the radius, right? not a 3d space?

It's interesting how this is so considering protons are not stationary.

i guess this answers my question. hard to believe though!

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

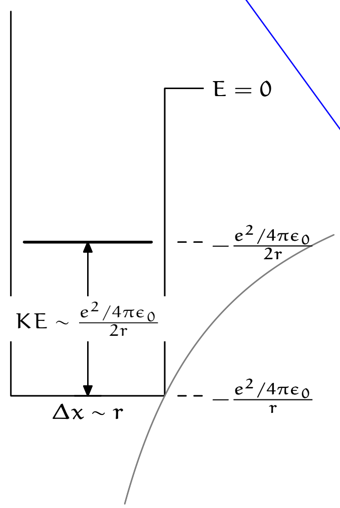

This energy is the confinement energy or the uncertainty energy.

This I think is actually directly contrary to what I learned in high school. Interesting!

I agree it's amazing how much we've been learning in this class that seems to contradict things we may have learned earlier on. One of the reasons I enjoy it so much

and i've seen at least 5 topics pop up in my other subjects!

these paragraphs are so hard to read when there's a hundred grey lines overlapping

stationary in respect to the center of the nucleus? because the electron is moving around right?

wait–what? do we really not understand this?

Possibly we either can't predict or explain it? At least not in the same way we can explain traditional mechanics

Part of the problem is that we don't have an intuitive feel for quantum mechanics. In fact, that's how quantum came about, sort of: physicists kept observing tings that were really weird and had to come up with new theories to explain them.

I think he says this because you can only explain qm mathematically and it's a little fuzzy.

I believe a lot of physicists in quantum mechanics have their own explanations/intuition on the matter, but there is no consistent 'correct' intuition on the matter, as there is with classical mechanics.

The reason he says that is because QM is really weird. "No one understands it" just means it violates every bit of our classical intuition possible.

I think I can safely say that nobody understands quantum mechanics. – Richard Feynman

And yet, we can still use the tools it provides, even if we don't exactly understand why.

I'm not sure why this is necessary... it seems like the approximation method should just be introduced rather than one nobody can use.

I think you could expand a little more from here?

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

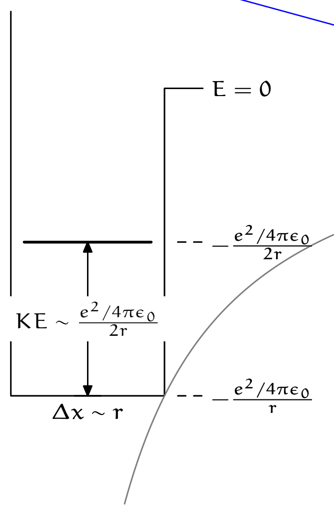

Having a mathematical model of some phenomenon confers at least SOME understanding, no?

that's an interesting statement. In my mind, a model is a representation of some physical process, but especially in the case of an a model of an unknown phenomenon, even if the model produces similar results - you could still be modeling the wrong thing. But I do agree, having some model tells you more than not having a model at all.

I think the disparity exists between understanding and explanation. A model can explain the behavior, but does not necessarily lead to a full understanding of why, and what individual aspects mean. Rather, it just tells you the end result for a given set of inputs.

When you say only understanding mathematically? Wouldn't they use some sort of approximation for this reasoning? Do you mean that the mathematical reasoning is not 100% accurate or trusted?

this makes me think about the steady state assumptions that we always make? does steady state really exist or is it just to make the problems simpler mathematicallly

At what scale? At the atomic scale, there really isn't steady state, but it isn't practical to look at lots of phenomena at that length scale. Almost any problem is going to involve some lumping.

Stationary state and steady state aren't the same thing, if that's the confusion.

is this paragraph necessary? I feel like its just a lot of background info into quantum.

I really liked this paragraph. It sets up the question well and allows you to visualize the problem. I would like if you did this on all of the problems we face.

Yes, background always helps. This section sure did spur a lot of debate, though. And clarify to me just how much I don't understand about the atomic world....

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

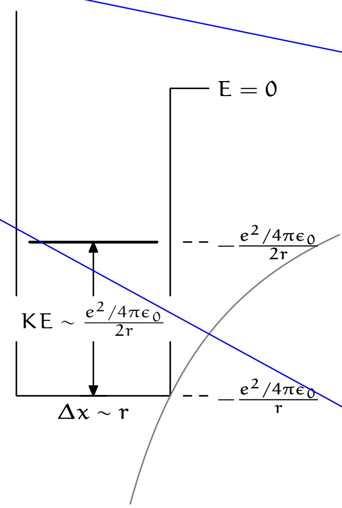

I remember in chemistry once my teacher talking about how we can model electrons as being stationary but wasn't explained like this and didn't make sense at the time. I have a much better understanding of the rationale now.

I don't really, but is this a form of lumping, to assume that it is stationary?

Correct, they are actually not only not stationary, electrons don't even follow circular orbits. In fact, the location of an electron cannot be predicted without losing some level of precision and information. Thus, we can lump this information and ignore such details.

Stationary in this sense literally means that the expectation value of the position of the electron does not depend on time. This falls out of quantum mechanics, it doesn't require lumping.

I am really having trouble understanding the electron being stationary. Is this a literal or figurative definition?

My understanding is that it's the latter, and I think a lot of people are getting tripped up on this issue in the other comments.

So is this the lumping right here? taking an electron to be "stationary" when we know that it actually can be anywhere in the universe at any given time?

Seems like it - the lumping is only paying attention to the probable location and then reducing that to a stationary location.

this was a funny word for me, I typically do consider stationary things to be able to radiate. When I was reading stationary I was just thinking, "no kinetic energy." (obviously that is not what he meant though)

I wonder exactly then how it does contribute KE?

What do you mean here by radiate? It doesn't radiate what?

Energy. He's saying here, that we assume it doesn't lose energy because it's stationary, so it doesn't radiate away energy and fall into the nucleus.

basically it's not using its KE.

I am still confused about this explanation of an electron

This is a lot of information to digest in one go.

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

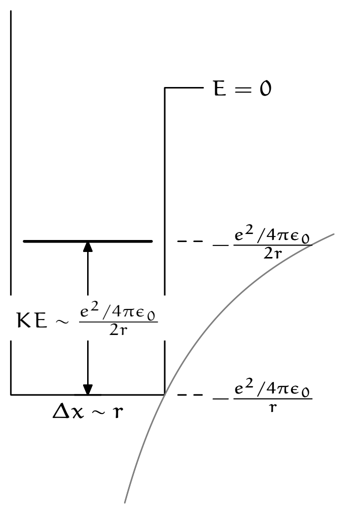

This energy is the confinement energy or the uncertainty energy.

this seems contradictory to what you just said...if the electron is stationary, shouldn't it exist as a single point? I think you should clarify your definition of radius here...

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

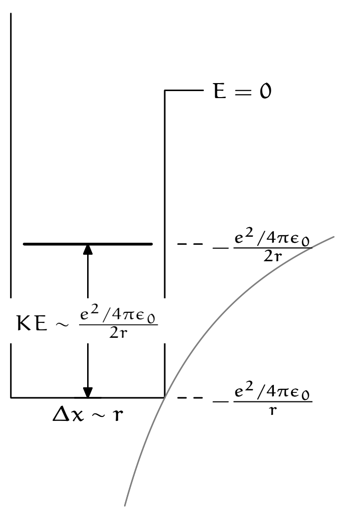

In that the probability of it being anywhere in the whole universe is non zero?

Yeah I think that's a much better way of putting it - I imagine something literally smeared all over the universe with Sanjoy's wording, like paint over a window or a bug on a windshield.

Which theory does this refer to? This is an amusing way to put it though.

haha I agree. It's a different sort of imagery.

Yeah, so I understand that in order to truly understand atoms, you need to include quantum mechanics. But I'm assuming this textbook is going to be used in other schools, where I'm guessing not all students are familiar with quantum mechanics. I just feel like this example is a bit complicated, or at least, some of the explanations are a bit complicated and unnecessary.

The same was true (and maybe moreso) for the earlier Bohr radius explanation. If this is intended for MIT students I think it's ok, but for younger students or ones without calculus, physics, computer science, or chemistry, there might be a problem. That's a running theme, though, which I think Sanjoy has thought about!

Personally I love these explanations no matter how complicated they get. I love how this class gives us the true explanations of things and not an oversimplified model which may hold many untruths. If anything these explanations make me look into these topics even more, which I would definitely not be doing otherwise.

I actually think this stuff was covered in my AP Chem class in high school. I think what he's talking about refers to the probability density that an electron's position around the nucleus. As I recall from my chem book, this appeared as a cloud around the nucleus when high probability at a certain band. By him saying that its actually smeared over the whole universe, I think he's saying that theoretically the electron could be anywhere just with smaller probability.

Yes, I really also love that these explanations, no matter how approximate they are, are based on intuitions of truths of physics. This brings to mind the paper that Sanjoy wrote about teaching 1st year physics in college I read.

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

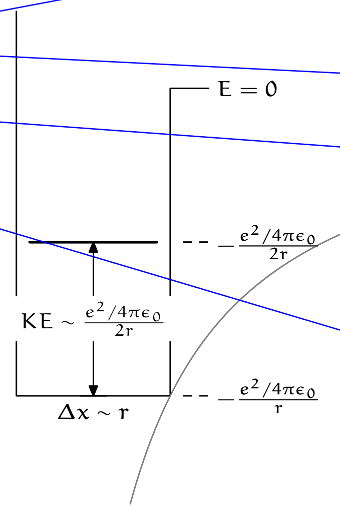

I want Sanjoy to sit in on some of my other classes and give those teachers some pointers!

Agreed.

This phrasing is a little confusing... I'm assuming this means the probability of it is more likely to lie in a typical radius, but the phrasing makes it sounds like this giant blob, a large portion of which lives in a confined area.

haha, I hadn't read it like that but I see what you mean. Would it make it better to say that for a significant amount of the time it lives within a typical radius?

I don't think that significant amount of time is appropriate, just because then it literally implies time. Even though it may live in the radius most of the time, I think it's important that people realize it's with a high probability...

I always wanted to learn this stuff, but never had the time to take an MIT physics class.

I feel like a few classes in course 6 (on the EE side) have teased about this but never really got into it - I'm glad we're looking at it in this one!

Can we get a drawing of this to show what a0 is? Or are r and a0 the same thing?

r is just the variable that we'll be using to eventually approximate a0.

We are trying to show that r=a0.

I like this idea! I'm excited for the section to follow..

How does confinement give energy? Is there no energy if the electron is not confined to a region?

Maybe energy kept because it IS confined? I dunno

It has to do with the allowable energy states for an electron and how the angular momentum has to be an integer multiple of Planck's constant.

What is confinement energy?

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

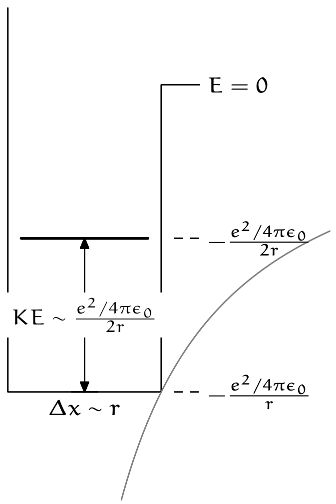

Insert "Heisenberg" here?

Agreed.

Are there other uncertainty principles?

credit where credit's due!

is this something that must be known before attempting this problem?

I understand this because I already took 8.04, but I think a sentence or 2 explaining what this equation means would be good so people could better understand what you are explaining

Yeah, I could use more explaining on this equation. how do we estimate uncertainty? are we simply relying on the picture and equations to the right?

Are the units of delta_x and delta_p in terms of distance and momentum? If feel like a larger momentum would result in a larger uncertainty in x, not lower. I might be thinking about this wrong.

Yea it would be helpful if there was a chart to follow on the side to help organize what is going on in this equation and the diagram.

can we go over this in class? i read it over a couple times and still don't get what "h" is

yeah I've never really learned this stuff, so a little more explanation of this before having the equation would be useful.

I think even before coming to MIT with very little physics background, I've heard of the uncertainty principle...

The reason he probably didn't put charts for the uncertainty principle, it's because it takes insane amounts of calculus and calculations in order to explain its origin. It is best for him to just mention it

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

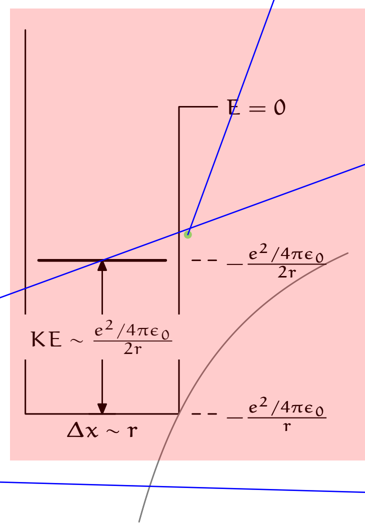

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

This graph (to me) is a little more messy than it could be. Labels and maybe making it slightly bigger might help.

and a little confusing- what exactly is it showing?

I agree. It took me a while to understand the picture.

To be honest, I'm still not sure what's going on here.

I understand (I think) the difference in energy due to the shells, but what is the (random?) arc on the right?

I agree, I had problems understanding the diagram since it was neither completely explained nor are there labels for everything.

I too am a little confused here. It would be nice if this graph had the same kind of breakdown as on the next page for E vs r.

It would also help to mention this figure at some point in the text. I am not familiar with quantum mechanics and don't fully understand what's going on in this graph.

I also don't know what the arc is, but could it be the orbit of highest probability?

What is going on in this graph? I tried to make sense of it but ended up having to ignore it instead.

agreed on the confusion of this picture

I assume the intermediate step here starts with E mvˆ2?

Yeah the equations aren't immediately obvious from the description in the text.

I was a little confused here too and had to stop reading and think about it for a bit. It might be helpful to slip in an equation or two.

And p = mv

Yeah, it would definitely be helpful here to mention that E m*(vˆ2), and also to mention that p = mv. Thank you to the person above who noted that p = mv, otherwise I still would have been lost. I'd definitely make these additions to the text, it'll flow nicer.

I think a better way to state this is that the rms value (?) for the momentum is approximately delta p. Then KE = pˆ2/2m.

7.3 Quantum mechanics: Hydrogen revisited

As a second computation of the Bohr radius a_0, here is a lumping and easy-cases method. The Bohr radius is the radius of the orbit with the lowest energy (the ground state). The energy is a sum of kinetic and potential energy. This division suggests, again, a divide-and-conquer approach: first the kinetic energy, then the potential energy.

What is the origin of the kinetic energy? The electron does not orbit in any classical sense. If it orbited, it would, as an accelerating charge, radiate energy and spiral into the nucleus. According to quantum mechanics, however, the proton confines the electron to a region of size r – still unknown to us – and the electron exists in a so-called stationary state. The nature of a stationary state is mysterious; no one understands quantum mechanics, so no one understands stationary states except mathematically. However, in an approximate estimate you can ignore details such as the meaning of a stationary state. The necessary information here is that the electron is, as the name of the state suggests, stationary: It does not radiate. The problem then is to find the size of the region to which the electron is confined. In reality the electron is smeared over the whole universe; however, a significant amount of it lives within a typical radius. This typical radius we estimate and call a_0.

For now let this radius be an unknown r and study how the kinetic energy depends on r. Confinement gives energy to the electron according to the uncertainty principle:

$$\Delta x \Delta p \sim \hbar,$$

where Δx is the position uncertainty and Δp is the momentum uncertainty of the electron. In this model $\Delta x \sim r$, as shown in the figure, so $\Delta p \sim \hbar/r$. The kinetic energy of the electron is

$$E_{\text{kinetic}} \sim \frac{(\Delta p)^2}{m_e} \sim \frac{\hbar^2}{m_e r^2}.$$

This energy is the confinement energy or the uncertainty energy.

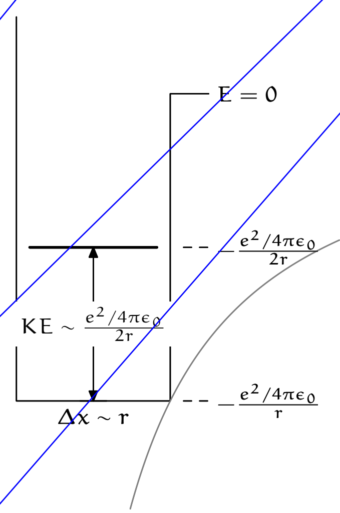

Should this be a delta E, if we're using delta p? (to signify the range of kinetic energy variations)

Why do we know that the 'confinement' energy is kinetic energy, as opposed to potential energy?

It comes out of the uncertainty relationship given above – "confinement" means restricting delta x, so delta p must be large enough that the product of those two terms is at least h. Delta p just means the distribution of momentum, so increasing it translates to increasing potential energy (which is, of course, due to momentum).

That's a good question...I'd like to know the answer as well.

The energy derived uses momentum and is thus related to motion and kinetic energy. That word "confinement" has to do with the model of an atom. The specified energy confines the electron to the atom.

I feel like confinement just kind of gives the idea of not kinetic... I guess that is not really a helpful explanation

This section had me reading all the peer notes I could, hoping for some clarifications that I couldn't find in the readings...

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small } r) \\ -1/r & \text{(large } r) \end{cases} \tag{7.5}$$

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

This is sort of a strange way to introduce the use of lumping in this section on lumping. It almost treats lumping as a afterthought to the example. I like (but don't entirely understand) this example, but I would like to see more emphasis on lumping.

I like it when the text states how the current problem is being solved by the overall section's new method. It's a good reminder and helps relate the problem to the approximation method.

I agree, it allows me pay attention to what he is doing as opposed to trying to figure out what method he is using.

maybe reference the lumping section here (i.e. (1.2.3)) just for people who skip around sections and might not know what this refers to

I'm interested to see how using lumping twice affects the accuracy of the results.

It seems like when we make multiple approximations, they are frequently tailored such that they will, to some extent, cancel each other out. It'll be interesting to see what happens here.

There's some sort of a trade-off. The less you lump, the fewer places to introduce error. The more you approximate, the better the chances that the errors cancel out.

I've never completely accepted this - although I can see how in _some_ estimations this would happen, I hardly think we can take that as an absolute rule.

when did we introduce this well?

The idea of "infinity" in physics and engineering always slightly bothered me because I can't really picture something with "infinite" anything.

I think a picture would really help here. I learned about this briefly in 6.007 but I still don't really understand it. Pictures help though, especially for such a theoretical subject.

Is this sort of represented in the diagram on the previous page? There are high sides to a well in that picture. On the difficulty picturing anything infinite, I think this is like a step function instead of one that's more curved. A step function has infinite slope but you can picture it.

I think he just means the electron is stuck in this shell. It isn't going anywhere. So it is in constant potential well with no chance of escaping.

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small r)} \\ -1/r & \text{(large r)} \end{cases} \tag{7.5}$$

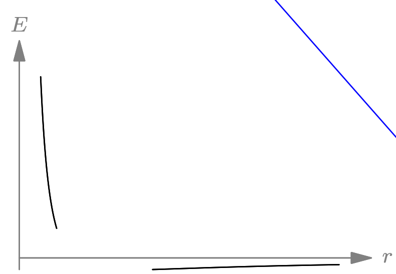

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

Based on this which things are still left uncertain. Just velocity?

What do you mean by high sides?

it means, that there is zero probability for it to exist outside the well

This paragraph makes how we used lumping very clear. I was initially confused, but these two sentences helped a lot.

Is there a better way to explain an "electron being smeared all over"? I understand what you're saying, but I'm guessing quite a few people in this class don't understand, and therefore plenty of people elsewhere reading this text would also be confused.

Maybe just "equal probability of being anywhere within the radial distance" or even "an even distribution of possible locations" or something

I think smeared is probably the easiest way to describe the phenomenon as far as readability goes for an average reader.

I'd personally find some comment regarding the fact that there's a non-zero probability of it existing in various places, but significantly greater probability at "r" (although this may not be quite right.)

Perhaps use "spread"?

I think you should mention that this is a probability density thing. It's most likely at radius r, so that's why we assume it, but it really could be anywhere.

I think you should mention that this is a probability density thing. It's most likely at radius r, so that's why we assume it, but it really could be anywhere.

Perfect, answered my earlier question.

Yep, good example of lumping approxmation.

I feel like it would be a little more helpful and clearer if this was mentioned when the approx. was made.

I really dislike this phrasing.

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small } r) \\ -1/r & \text{(large } r) \end{cases} \tag{7.5}$$

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

I thought the first part of our lumping was to replace the electrostatic potential with potential well... but the entire problem seems to use the electrostatic potential and not a potential well?

Yeah, I am confused by this as well.

I don't remember why this is negative.

I believe its because of the direction at which you define your E field.

I think it has something to do with an electron having a negative charge.

The general equation has q1*q2 in the numerator, so in this case we have the 2 charges of the electron (-e) and the proton (e), so we get -eˆ2. Two oppositely charged particles very far apart then have a potential energy approaching 0 and as they get closer, it gets more negative, sort of the reverse of gravity. If they were attracting each other starting from very far apart, their original potential energy (zero) would be converted into KE as they accelerated toward each other, thus LOWERING PE and making it negative.

also, this is so that likes repel and opposites attract.

Because the potential energy is attractive (i.e. between an electron and a proton). It would be positive if the potential energy was repellent.

Basically, it's negative if the charges have opposite sign and positive if the charges have the same sign.

What's wrong with taking the derivative of the energy to find the value of r that minimizes the energy? I don't think it's overly-complex using calculus to find a First Order Condition that yields the r we are looking for. Plus, it eliminates any errors we might have made in figuring out what term dominates at large and small r's.

Because this section is all about anti-calculus! Wouldn't it be more fun to approximate it?

but it's really easy to find the max of A/rˆa + B/rˆb=C. It's just rˆ(b-a)=-(b*B)/(a*A). (A handy abstraction for you.)

Here a=1, b=2, so r=2*(hˆ2/m)/(eˆ2/4.pi.eps)

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small } r) \\ -1/r & \text{(large } r) \end{cases} \tag{7.5}$$

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

This is a familiar technique for sketching out curves. I never thought of it as it's own approximation technique however.

Yeah, this is the whole idea of easy cases–to take a complicated equation or problem and apply it to some simple cases that may cause several variables to drop out, etc.

I really like how this section is pulling out a lot of old techniques like divide and conquer and easy cases.

I like how this example uses both the "easy cases" method as well as lumping. It helps tie the two different methods together and reinforce the stuff we learned earlier.

I like this simple explanation.

Yeah, it was super easy for me to follow and understand despite my lack of knowledge in this area.

is this the lumping technique? simplifying all cases as one of these two?

This almost feels like a sub-step using easy cases, and then we use this to lump results together. Just my guess.

How do we have any idea what the scale of r is, does this change the calculations at all?

I really like how the readings apply previous concepts to solving new and more complicated problems as we go along. It really helps to demonstrate the power these techniques have when used in conjunction with eachother.

Again, I like the simple breakdown and repeat

I agree, this made what could have been a very complicated section much easier to follow.

Is this an example of how the math explains the physics, i.e. stating how we don't understand quantum physics?

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{7.5}$$

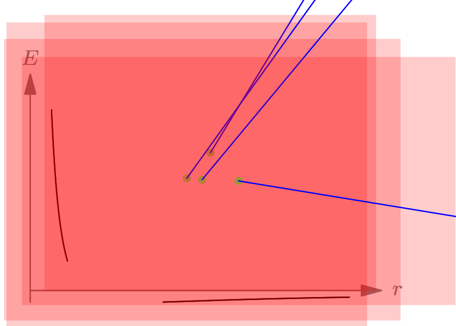

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

I figured that we always do this in engineering

Yeah, lumping is used really often in engineering and even physics.

I agree but I think easy cases is the more applicable technique. I have always found that I have been able to solve problems that others can't simply by looking at the easy cases if I don't know the formula necessary for finding the exact answer.

agreed that we do this in engineering all the time though, is 1st(2nd?) order taylor approximation necessary for this since there i some degree of r in the denominator in both terms?

This graph and explanation behind it are really cool. The graph is really similar to those created using wavefunctions in chemistry. I'm assuming this is no coincidence.

I remember talking about this in chemistry and 2.002

At first glance this could almost seem like random marks on the graph, but thinking about it more helps me understand why the extreme cases can help us figure out what goes on in the middle.

I agree. Based on the two lines here, my intuition tells me that there is really only one *likely* solution (granted, it could do all kinds of crazy stuff in the middle... but let's hope not...) - which was confirmed when I looked at the final sketch.

I think this graph is really helpful. since we know that initially the graph is decreasing and for large r, the graph is increasing, there must be a minimum somewhere in between if the graph is continuous. we know it's continuous because the value we're looking for here is Energy (kinetic+potential), and energy is conserved, without any external forces

Yeah this is a good example of putting down what you know and then kind of connecting the dots. It works well.

why wasn't this done on a log-log plot? Theyre quite useful, and it seems like this course is trying to emphasize them somewhat. This seems like a prime opportunity outside of the realm of bode plots.

Log-log plots can't handle negative numbers.

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small r)} \\ -1/r & \text{(large r)} \end{cases} \quad (7.5)$$

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

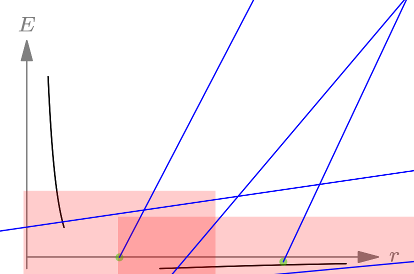

maybe emphasize the importance of the energy value and corresponding radius?

I don't think there is any need to – this makes sense as is, especially to someone who knows little about these things. It's interesting from an analytic perspective.

What if the line is horizontal? That doesn't count as minimum, does it?

I'm not positive it would apply here, but it would be nice to see a dotted line connecting the two cases on this graph

But then we wouldn't have our own mind approximate what might happen in between the end behavior and I think that's the point here.

Yes, maybe more of a point-out in the text to the missing chunk then.

why is this line almost horizontal, if it had a slightly negative slope, then the function wouldn't have a minimum except for r = infinity, how do we know that this isn't the case from the drawing?

To me it looks like the slope dE/dr is 1/r^2 for small values and -1/r^3 for large values.

Meaning that the slope at small r is positive and negative for large r, opposite of what the text says.

Someone please point out if I'm making a stupid mistake

take a really close look at the graphs - what the text says is correct.

The pictures are E(r) functions. For small values of r, the value of E decrease as r increases, so there would be a negative slope. For large values of r, E would increase as r increases, so there's a positive slope.

You can also see this by taking the derivative of 1/r^2 and -1/r.

Assuming the derivative is monotonic increasing...

Maybe give a brief note about how large r physically makes sense to have increasing energy and smaller r values are associated with decreasing energy... It makes sense on the graph though.

I like this reasoning

This estimate uses lumping twice. First, the complicated electrostatic potential, which varies with distance, is replaced by the simple potential well (with infinitely high sides). Second, the electron, which in reality is smeared all over, is assumed to be at only one spot that is at a distance r from the proton.

This second lumping approximation also helps us estimate the potential energy. It is the classical electrostatic energy of a proton and electron separated by r:

$$E_{\text{potential}} \sim -\frac{e^2}{4\pi\epsilon_0 r}.$$

Therefore, the total energy is the sum

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}.$$

With this total energy as its guide, the electron adjusts its separation r to make the energy a minimum. As a step toward finding that separation, sketch the energy. For sketching functions, the first tool to try is easy cases. Here, the easy cases are small and large r. At small r, kinetic energy is the important term because its $1/r^2$ overwhelms the $1/r$ in the potential energy. At large r, potential energy is the important term, because its $1/r$ goes to zero more slowly than the $1/r^2$ in the kinetic energy.

$$E \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{7.5}$$

Now we can sketch E in the two extreme cases. The sketch demonstrates the result in which we are interested: that there must be a minimum combined energy at an intermediate value of r. There is no smooth way to connect the two extreme segments without introducing a minimum. An analytic argument confirms that pictorial reasoning. At small r, the slope dE/dr is negative. At large r, it is positive. At an intermediate r, the slope must cross between positive and negative. In other words, somewhere in them middle the slope must be zero, and the energy must therefore be a minimum.

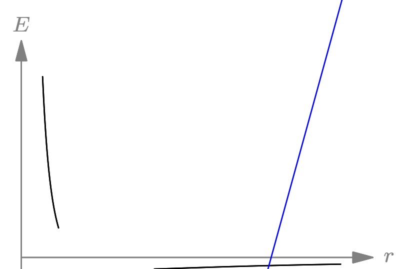

typo

I really appreciate the detailed walkthrough of the graph and it's meaning. It helped to slow things down and allow me to digest the problem so far.

Yeah, I never thought of thinking about deriving the wave functions in this way.

There are two approximate methods to determine the minimum r. The first method is familiar from the analysis of lift in Section 3.5.2: When two terms compete, the minimum occurs when the terms are roughly equal. In other words, at the minimum energy, the potential energy and kinetic energy (the two competing terms) are roughly equal in magnitude. Using the Bohr radius a_0 as the corresponding separation, this criterion says

$$\underbrace{\frac{e^2}{4\pi\epsilon_0 a_0}}_{\text{PE}} \sim \underbrace{\frac{\hbar^2}{m_e a_0^2}}_{\text{KE}}. \quad (7.6)$$

The result for a_0 is

$$a_0 \sim \frac{\hbar^2}{m_e(e^2/4/4\pi\epsilon_0)}. \quad (7.7)$$

The second method of estimating the minimum-energy separation is to use dimensional analysis, by writing the energy and radius in dimensionless forms. Such a rewriting is not mandatory in this example, but it is helpful in complicated examples and is therefore worth learning via this example. To make r dimensionless, cook up another length l and then define $\bar{r} \equiv r/l$. The only other length that is based on the parameters of hydrogen (and the relevant constants of nature such as ϵ_0) is

$$l \equiv \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)}.$$

So, let's defined the scaled (dimensionless) radius as

$$\bar{r} \equiv \frac{r}{l}.$$

To make the energy dimensionless, cook up another energy based on the parameters of the problem. A reasonable candidate for this energy scale is $e^2/4\pi\epsilon_0 l$. That choice defines the scaled energy as

$$\bar{E} \equiv \frac{E}{e^2/4\pi\epsilon_0 l}.$$

Using the scaled length and energy, the total energy

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}$$

simplifies greatly:

should this be "approximation methods"?

Does this mean for two terms that equivalent to something? ie. two terms that are added, multipled, together

I need to remember this!

agreed- i wouldn't have known that otherwise.

This reminds me of the minimization/maximization problems from middle and high school where we had to graph out the functions constraining us to solve the problem visually, and always the answer was at the intersection of the functions.

Reminds me of the drag problem from way back when about the ideal velocity or height for a plane to fly.

I'm pretty sure that's the exact example he's talking about

This is pretty intuitive.

Yea, couldn't we just observe the derivative graph and determine the minimum?

I'd recommend against memorizing this. It's wrong when the powers aren't equal in magnitude and opposite in sign. It's tempting to use this sort of trick more generally, but it can easily put you off by a factor of few.

I wrote this elsewhere, but I find it much more useful: the max of A/r^a+B/r^b=C is at r^(b-a)=-(b*B)/(a*A). It differs by a factor of b/a from assuming the two terms are equal, but it's exactly right and applicable everywhere, so it's a much more useful abstraction.

I'm a bit confused and not quite sure what this means?

are you saying this is true at the minimum r?

There are two approximate methods to determine the minimum r. The first method is familiar from the analysis of lift in Section 3.5.2: When two terms compete, the minimum occurs when the terms are roughly equal. In other words, at the minimum energy, the potential energy and kinetic energy (the two competing terms) are roughly equal in magnitude. Using the Bohr radius a_0 as the corresponding separation, this criterion says

$$\underbrace{\frac{e^2}{4\pi\epsilon_0 a_0}}_{\text{PE}} \sim \underbrace{\frac{\hbar^2}{m_e a_0^2}}_{\text{KE}}. \tag{7.6}$$

The result for a_0 is

$$a_0 \sim \frac{\hbar^2}{m_e(e^2/4/4\pi\epsilon_0)}. \tag{7.7}$$

The second method of estimating the minimum-energy separation is to use dimensional analysis, by writing the energy and radius in dimensionless forms. Such a rewriting is not mandatory in this example, but it is helpful in complicated examples and is therefore worth learning via this example. To make r dimensionless, cook up another length l and then define $\bar{r} \equiv r/l$. The only other length that is based on the parameters of hydrogen (and the relevant constants of nature such as ϵ_0) is

$$l \equiv \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)}.$$

So, let's defined the scaled (dimensionless) radius as

$$\bar{r} \equiv \frac{r}{l}.$$

To make the energy dimensionless, cook up another energy based on the parameters of the problem. A reasonable candidate for this energy scale is $e^2/4\pi\epsilon_0 l$. That choice defines the scaled energy as

$$\bar{E} \equiv \frac{E}{e^2/4\pi\epsilon_0 l}.$$

Using the scaled length and energy, the total energy

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}$$

simplifies greatly:

Does it matter that we just dropped the negative here?

I doubt it, since we're just using twiddle and not equals.

He's just setting the magnitude of the energies equal here. I don't think the sign matters in this situation.

with a twiddle, "dropping" the negative would be like making the constant negative.

I love how twiddle is now acceptable in mathematical discussion.

i don't fully understand where the lumping came in

An extraneous /4 here?

Ooh good catch. I totally missed it on my pass through of the equation.

I think so

Yeah, pretty sure.

I was actually just thinking to myself if we were going to do this!

can you do this?

I think he just wants a unitless value...but right now it seems sketchy

What exactly does "L" represent?

I don't remember ever cooking anything up in the dimension section. Why is this necessary, how is it legit, and how do we know when to apply this?

I think he did this because most equations become easier to work with when you take out their dimensions. He cooked up this value because it is a length specified only by the quantities relevant to the problem and therefore defines some sort of length scale the problem has.

There are two approximate methods to determine the minimum r. The first method is familiar from the analysis of lift in Section 3.5.2: When two terms compete, the minimum occurs when the terms are roughly equal. In other words, at the minimum energy, the potential energy and kinetic energy (the two competing terms) are roughly equal in magnitude. Using the Bohr radius a_0 as the corresponding separation, this criterion says

$$\underbrace{\frac{e^2}{4\pi\epsilon_0 a_0}}_{\text{PE}} \sim \underbrace{\frac{\hbar^2}{m_e a_0^2}}_{\text{KE}}. \tag{7.6}$$

The result for a_0 is

$$a_0 \sim \frac{\hbar^2}{m_e(e^2/4/4\pi\epsilon_0)}. \tag{7.7}$$

The second method of estimating the minimum-energy separation is to use dimensional analysis, by writing the energy and radius in dimensionless forms. Such a rewriting is not mandatory in this example, but it is helpful in complicated examples and is therefore worth learning via this example. To make r dimensionless, cook up another length l and then define $\bar{r} \equiv r/l$. The only other length that is based on the parameters of hydrogen (and the relevant constants of nature such as ϵ_0) is

$$l \equiv \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)}.$$

So, let's defined the scaled (dimensionless) radius as

$$\bar{r} \equiv \frac{r}{l}.$$

To make the energy dimensionless, cook up another energy based on the parameters of the problem. A reasonable candidate for this energy scale is $e^2/4\pi\epsilon_0 l$. That choice defines the scaled energy as

$$\bar{E} \equiv \frac{E}{e^2/4\pi\epsilon_0 l}.$$

Using the scaled length and energy, the total energy

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}$$

simplifies greatly:

I am confused... in dimensional analysis we looked at the variables we already have and made dimensionless groups using those.. why are we now making up new variables? it seems like it's going to complicate things and is also not an intuitive thought...

We're nondimensionalizing. It's one of the common techniques to make equations and variables simple:

http://en.wikipedia.org/wiki/Nondimensionalization

nondimensionalizing equations is one of the best skills I've attained at MIT so far.

Would this section be better served by being moved to the dimensional analysis section of the book? Then, you can just refer to it here. I understand the importance of doing things multiple ways, but the switching back and forth between methods in a section devoted to lumping makes the entire section a bit disjointed. This might just be a personal thing though.

I think it's fine here...I don't think it's great to limit a certain technique to a certain section and only focus on that technique because all of these techniques are meant to complement the others/double check yourself.

Is l just a0 * 4? So confused as to what l represents...

I don't understand where this came from. I thought we were using a new method to derive this. Or are we just using this result from above?

I think this is the "cooked-up" l? However, I don't really understand how it is chosen...

Perhaps he used dimensional analysis to get something with the dimension of length, and it happened right off to be this? (He'd have to have done the assumed grouping from a previous reading to get the constants in there.)

I believe he is using a length which is on the same magnitude as the radius and more importantly one that depends on hydrogen, which we are looking to solve for.

I think you are right. He probably used dimensional analysis and guessed the other variables that were important and used them to create a dimensionless group with l.

i would have never been able to come up with this – it's slightly frustrating in that sense

There are two approximate methods to determine the minimum r. The first method is familiar from the analysis of lift in Section 3.5.2: When two terms compete, the minimum occurs when the terms are roughly equal. In other words, at the minimum energy, the potential energy and kinetic energy (the two competing terms) are roughly equal in magnitude. Using the Bohr radius a_0 as the corresponding separation, this criterion says

$$\underbrace{\frac{e^2}{4\pi\epsilon_0 a_0}}_{\text{PE}} \sim \underbrace{\frac{\hbar^2}{m_e a_0^2}}_{\text{KE}}. \tag{7.6}$$

The result for a_0 is

$$a_0 \sim \frac{\hbar^2}{m_e(e^2/4/4\pi\epsilon_0)}. \tag{7.7}$$

The second method of estimating the minimum-energy separation is to use dimensional analysis, by writing the energy and radius in dimensionless forms. Such a rewriting is not mandatory in this example, but it is helpful in complicated examples and is therefore worth learning via this example. To make r dimensionless, cook up another length l and then define $\bar{r} \equiv r/l$. The only other length that is based on the parameters of hydrogen (and the relevant constants of nature such as ϵ_0) is

$$l \equiv \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)}.$$

So, let's defined the scaled (dimensionless) radius as

$$\bar{r} \equiv \frac{r}{l}.$$

To make the energy dimensionless, cook up another energy based on the parameters of the problem. A reasonable candidate for this energy scale is $e^2/4\pi\epsilon_0 l$. That choice defines the scaled energy as

$$\bar{E} \equiv \frac{E}{e^2/4\pi\epsilon_0 l}.$$

Using the scaled length and energy, the total energy

$$E = E_{\text{potential}} + E_{\text{kinetic}} \sim -\frac{e^2}{4\pi\epsilon_0 r} + \frac{\hbar^2}{m_e r^2}$$

simplifies greatly:

typo

How would I know what is reasonable or not?

Why is this the case, How are these parameters chosen?

Can you maybe explain here why it is a reasonable candidate?

despite the fact that these are familiar terms and equations, i feel like estimation problems are more fun with tangible figures and values.

$$\bar{E} \sim -\frac{1}{\bar{r}} + \frac{1}{\bar{r}^2}.$$

The ugly constants all live in the definitions of scaled length and energy. This dimensionless energy is easy to think about and to sketch.

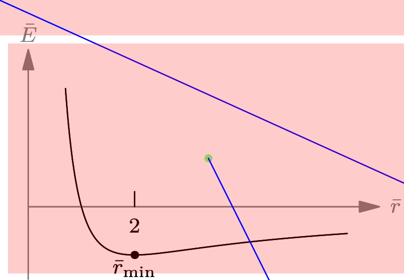

Calculus locates this minimum-energy $\bar{r}$ at $\bar{r}_{min} = 2$. Equating the two terms $\bar{r}^{-1}$ and $\bar{r}^{-2}$ gives $\bar{r}_{min} \sim 1$. In normal, unscaled terms, it is

$$r_{min} = l\bar{r}_{min} = \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)},$$

which is the Bohr radius as computed using dimensional analysis (Section 5.3.1) and is also the exact Bohr radius computed properly using quantum mechanics. The sloppiness in estimating the kinetic and potential energies has canceled the error introduced by cheap minimization! Even if the method were not so charmed, there is no point in doing a proper calculus minimization. Given the inaccuracies in the rest of the derivation, The calculus method is too accurate

Engineers understand this idea of not over-engineering a system. If a bicycle most often breaks at welds in the frame, there is little point replacing the metal between the welds with expensive, high-strength aerospace materials. The new materials might last 100 years, but such a replacement would be overengineering because something else will break before 100 years are done.

In estimating the Bohr radius, the kinetic-energy estimate uses a crude form of the uncertainty principle, $\Delta p \Delta x \sim \hbar$, whereas the true statement is that $\Delta p \Delta x \geqslant \hbar/2$. The estimate also uses the approximation $E_{kinetic} \sim (\Delta p)^2/m$. This approximation contains m instead of $2m$ in the denominator. It also assumes that Δp can be converted into an energy as though it were a true momentum rather than merely a crude estimate for the root-mean-square momentum. The potential- and kinetic-energy estimates use a crude definition of position uncertainty Δx: that $\Delta x \sim r$. After making so many approximations, it is pointless to minimize the result using the elephant gun of differential calculus. The approximate method is as accurate as, or perhaps more accurate than the approximations in the energy.

Comments on page 4

I dont quite get how you made this jump. Could you maybe show the simplification in class?

I'm actually surprised we waited as long as we did to reduce the equation to this. I feel as though we carried the constants for awhile when we weren't really using them.

Agreed. Didn't we come to this conclusion 2 pages ago in (7.5).

I think it was necessary though to show the entire process.

Yeah I found seeing the process drawn out to be helpful.

Lumping doesn't seem to be a technique that is useable on its own, most of this problem was still done with dimensional analysis- we're just naming the assumptions we've made all along as lumping

Well, many of the other techniques aren't necessarily useful "on their own" either. It depends on the problem and how you approach it. Often here, we've dealt with lumping and dimensional analysis or easy cases but that doesn't mean lumping is exclusively usable in conjunction with other skills.

Lumping is a way to simplify a complex problem to the sort that is more easily solvable, but you still need other methods to reach an answer.

I thought we were trying to avoid calculus? Wasn't lumping the "opposite" of calculus? Seems a little silly if we can't finish the problem without use of calculus ;)

Maybe because it's easy calculus?

I think he's just showing how 'close' our answer gets. Sort of like telling us the answer to book sales after we estimate it.

so is it 2 or 1 that we are gonna use?

Nice to see this graph again with the kinetic and potential energies connected

Great. Now it makes sense graphically.

can't beat graphical learning

$$\bar{E} \sim -\frac{1}{\bar{r}} + \frac{1}{\bar{r}^2}.$$

The ugly constants all live in the definitions of scaled length and energy. This dimensionless energy is easy to think about and to sketch.

Calculus locates this minimum-energy $\bar{r}$ at $\bar{r}_{min} = 2$. Equating the two terms $\bar{r}^{-1}$ and $\bar{r}^{-2}$ gives $\bar{r}_{min} \sim 1$. In normal, unscaled terms, it is

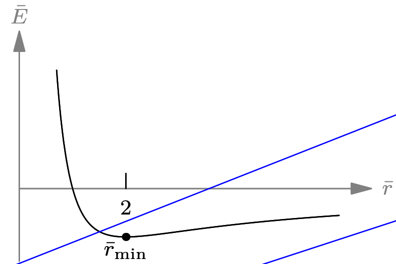

$$r_{min} = l\bar{r}_{min} = \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)},$$

which is the Bohr radius as computed using dimensional analysis (Section 5.3.1) and is also the exact Bohr radius computed properly using quantum mechanics. The sloppiness in estimating the kinetic and potential energies has canceled the error introduced by cheap minimization! Even if the method were not so charmed, there is no point in doing a proper calculus minimization. Given the inaccuracies in the rest of the derivation, The calculus method is too accurate

Engineers understand this idea of not over-engineering a system. If a bicycle most often breaks at welds in the frame, there is little point replacing the metal between the welds with expensive, high-strength aerospace materials. The new materials might last 100 years, but such a replacement would be overengineering because something else will break before 100 years are done.

In estimating the Bohr radius, the kinetic-energy estimate uses a crude form of the uncertainty principle, $\Delta p \Delta x \sim \hbar$, whereas the true statement is that $\Delta p \Delta x \geqslant \hbar/2$. The estimate also uses the approximation $E_{kinetic} \sim (\Delta p)^2/m$. This approximation contains m instead of $2m$ in the denominator. It also assumes that Δp can be converted into an energy as though it were a true momentum rather than merely a crude estimate for the root-mean-square momentum. The potential- and kinetic-energy estimates use a crude definition of position uncertainty Δx: that $\Delta x \sim r$. After making so many approximations, it is pointless to minimize the result using the elephant gun of differential calculus. The approximate method is as accurate as, or perhaps more accurate than the approximations in the energy.

somehow i like this method more because it seems more accurate because we are actually doing derivatives

Messy ones, but I see the point.

Is there some reason why this happened other than pure luck? Is the lesson that you are trying to tell us that we can be sloppy and it all will probably just work out, or was this a special case?

I don't think I understanding this. Are you saying that since we're only looking for the minimum, the sloppiness sort of doesn't matter?

What does this refer too?

equating the two energies to get an approximate minimum

This should be a lower case t.

There's a missing period here.

due to the fact that in quantum there are discrete levels of energy, (not continuous).

I'm not sure that's what he means. I think he might just mean a proper calculus method is like carrying 10 significant figures, when all of the rest of derivation only has accuracy to 2 significant figures.

But why does that make the answer worse?

It doesn't make the answer worse so much as it adds specificity where there was none to begin with.

it's similar to what we learn with sig figs, "you can't be more precise than your measuring device" – if we're going to estimate for the rest of the derivation, no point in being precise for this one part

Oh, that makes more sense.

$$\bar{E} \sim -\frac{1}{\bar{r}} + \frac{1}{\bar{r}^2}.$$

The ugly constants all live in the definitions of scaled length and energy. This dimensionless energy is easy to think about and to sketch.

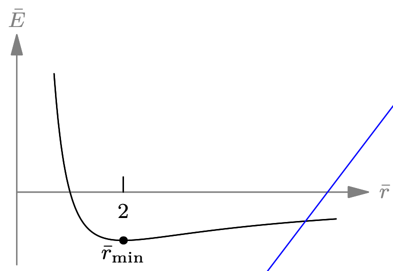

Calculus locates this minimum-energy $\bar{r}$ at $\bar{r}_{min} = 2$. Equating the two terms $\bar{r}^{-1}$ and $\bar{r}^{-2}$ gives $\bar{r}_{min} \sim 1$. In normal, unscaled terms, it is

$$r_{min} = l\bar{r}_{min} = \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)},$$

which is the Bohr radius as computed using dimensional analysis (Section 5.3.1) and is also the exact Bohr radius computed properly using quantum mechanics. The sloppiness in estimating the kinetic and potential energies has canceled the error introduced by cheap minimization! Even if the method were not so charmed, there is no point in doing a proper calculus minimization. Given the inaccuracies in the rest of the derivation, The calculus method is too accurate

Engineers understand this idea of not over-engineering a system. If a bicycle most often breaks at welds in the frame, there is little point replacing the metal between the welds with expensive, high-strength aerospace materials. The new materials might last 100 years, but such a replacement would be overengineering because something else will break before 100 years are done.

In estimating the Bohr radius, the kinetic-energy estimate uses a crude form of the uncertainty principle, $\Delta p \Delta x \sim \hbar$, whereas the true statement is that $\Delta p \Delta x \geqslant \hbar/2$. The estimate also uses the approximation $E_{kinetic} \sim (\Delta p)^2/m$. This approximation contains m instead of $2m$ in the denominator. It also assumes that Δp can be converted into an energy as though it were a true momentum rather than merely a crude estimate for the root-mean-square momentum. The potential- and kinetic-energy estimates use a crude definition of position uncertainty Δx: that $\Delta x \sim r$. After making so many approximations, it is pointless to minimize the result using the elephant gun of differential calculus. The approximate method is as accurate as, or perhaps more accurate than the approximations in the energy.

do you mean over complicating?

Over engineering could include over complicating but it's more broad than that. http://en.wikipedia

I think he's using this to mean the same thing as over-complicating, he uses the same phrase again later in the paragraph so i think he's just coining a new term :)

the example he gives of the bike sounds like overengineering has less to do with over complicating and more to do with needless redundancy and inefficiency.

Sounds like its just not balanced, if you make the whole bike with comparably good materials it should all last 100 years until the weld breaks but if you make the other parts with materials that will last as long as the estimated time to failure of the welds then you can be efficient.

It's basically the idea of if it's not the weakest link, don't fix it yet.

Huh? how is this relevant?

I'm not entirely sure it's exact relevance, but it is a cool and truthful anecdote!

I think this is following up on the "calculus method is too accurate" comment, meaning that we shouldn't "over-engineer" by using calculus when the rest of approximation is inaccurate. Likewise, we shouldn't use aerospace materials to patch a cheap bicycle.

That explanation helps a bit in understanding how not to do too many complicated calcutions. It is just hard to draw any parallels between the two examples.

I think this analogy is very fitting for the situation.

Yeah I like it.

Great example to show why it's bad to over-engineer a system. I guess I've always been a fan of using calculus or "over-engineering" problems in this course thus far, but this analogy sheds some light on not doing that.

Haha, yeah. Imagine the costs.

Agreed. I like how we are relating our choice of math/engineering to practical applications.

$$\bar{E} \sim -\frac{1}{\bar{r}} + \frac{1}{\bar{r}^2}.$$

The ugly constants all live in the definitions of scaled length and energy. This dimensionless energy is easy to think about and to sketch.

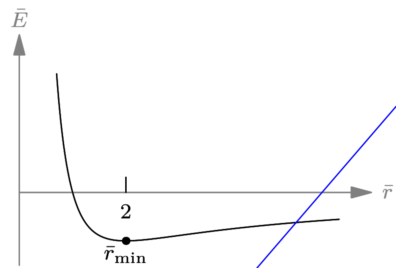

Calculus locates this minimum-energy $\bar{r}$ at $\bar{r}_{min} = 2$. Equating the two terms $\bar{r}^{-1}$ and $\bar{r}^{-2}$ gives $\bar{r}_{min} \sim 1$. In normal, unscaled terms, it is

$$r_{min} = l\bar{r}_{min} = \frac{\hbar^2}{m_e(e^2/4\pi\epsilon_0)},$$

which is the Bohr radius as computed using dimensional analysis (Section 5.3.1) and is also the exact Bohr radius computed properly using quantum mechanics. The sloppiness in estimating the kinetic and potential energies has canceled the error introduced by cheap minimization! Even if the method were not so charmed, there is no point in doing a proper calculus minimization. Given the inaccuracies in the rest of the derivation, The calculus method is too accurate

Engineers understand this idea of not over-engineering a system. If a bicycle most often breaks at welds in the frame, there is little point replacing the metal between the welds with expensive, high-strength aerospace materials. The new materials might last 100 years, but such a replacement would be overengineering because something else will break before 100 years are done.

In estimating the Bohr radius, the kinetic-energy estimate uses a crude form of the uncertainty principle, $\Delta p \Delta x \sim \hbar$, whereas the true statement is that $\Delta p \Delta x \geqslant \hbar/2$. The estimate also uses the approximation $E_{kinetic} \sim (\Delta p)^2/m$. This approximation contains m instead of $2m$ in the denominator. It also assumes that Δp can be converted into an energy as though it were a true momentum rather than merely a crude estimate for the root-mean-square momentum. The potential- and kinetic-energy estimates use a crude definition of position uncertainty Δx: that $\Delta x \sim r$. After making so many approximations, it is pointless to minimize the result using the elephant gun of differential calculus. The approximate method is as accurate as, or perhaps more accurate than the approximations in the energy.

Why did you bother to make this 'approximations' when the exact equation was known? I think for many people you already sort of pulled these equations out of a hat, so why not pull out the correct equations to begin with?

he explains below, that that ould be overengineering. We know the uncertainty principle so having some sort of approximation that deviates from its actual value by less than the uncertainty value, doesn't affect out answer at all.

I like the recap of everything you approximated.

I'm not sure what I'm supposed to be learning from these paragraphs – there is so much information and so many things I don't know.

tis is a very useful point to remeber. u should bold it to make sure readers get the point

I'm pretty surprised by this, it just seems like something that we would have already figured out the exact method for solving.

I wish this was gone over in more detail, with explanations as to why it is more accurate

Same here, perhaps a little more explanation on accuracy would help ease my unrest here.

The method of equating competing terms is called balancing. We balanced the kinetic energy against the potential energy by assuming that they are roughly the same size. Nature could have been unkind: The potential and kinetic energies could have differed by a factor of 10 or 100. But Nature is kind: The two energies are roughly equal, except for a constant that is nearly 1 (of order unity). This rough equality occurs in many examples: You often get a reasonable answer simply by pretending that two energies (or two quantities with the same units) are equal. [When the quantities are potential and kinetic energy, as they often are, you get extra safety: The so-called virial theorem protects you against large errors (for more on the virial theorem, see any intermediate textbook on classical dynamics).]

Yep. I find myself doing this all the time.

is this another method to think about? or just an explanation of above?

In general, will we be able to assume that KE and PE are the same size? Or should we assume that it only worked in this case?

Only in this case, I think

Yeah, I like this method but am still a bit unsure about how often I can use it and in what cases. It would be nice if this was a section in itself.

Nature seems to be kind to us more often than not when we do these estimations.

What happens when nature is unkind? ie. we can assume that it will be a factor of 10 or 100 difference?

I have the same question. How do you decide when they are approximately the same or not?

Is there some way of knowing if this is true in an individual case? Or throughout this class (and life beyond) do we just assume it to be true and cross our fingers?

I think the idea is to assume it's true and cross our fingers, because it usually is for the degree of accuracy we need. If it's off it's usually a quick fix of a constant, and at least we've already done most of the work.

I think this goes in the category of its better to make a model and get some sort of an answer then get stuck trying to be more exact and getting no answer at all.

just use 1 here, stick with your defined conventions

also, can you talk about the limitation of using balancing? when should you NOT use it

when the ratio of the magnitudes of the powers is too large of an error factor (since that is precisely how much the answer from balancing differs from the calculus answer).

it would be great if you can just briefly mention some examples, it just makes the arguments more convincing

it would be interesting to hear more about why these energies are equal

The method of equating competing terms is called balancing. We balanced the kinetic energy against the potential energy by assuming that they are roughly the same size. Nature could have been unkind: The potential and kinetic energies could have differed by a factor of 10 or 100. But Nature is kind: The two energies are roughly equal, except for a constant that is nearly 1 (of order unity). This rough equality occurs in many examples: You often get a reasonable answer simply by pretending that two energies (or two quantities with the same units) are equal. [When the quantities are potential and kinetic energy, as they often are, you get extra safety: The so-called virial theorem protects you against large errors (for more on the virial theorem, see any intermediate textbook on classical dynamics).]

it'd be nice to have a better idea of when it isn't appropriate to do this

http://en.wikipedia.org/wiki/Virial_theorem

Thanks, this was helpful

Thanks.

I was wondering if that was a typo.

The method of equating competing terms is called balancing. We balanced the kinetic energy against the potential energy by assuming that they are roughly the same size. Nature could have been unkind: The potential and kinetic energies could have differed by a factor of 10 or 100. But Nature is kind: The two energies are roughly equal, except for a constant that is nearly 1 (of order unity). This rough equality occurs in many examples: You often get a reasonable answer simply by pretending that two energies (or two quantities with the same units) are equal. [When the quantities are potential and kinetic energy, as they often are, you get extra safety: The so-called virial theorem protects you against large errors (for more on the virial theorem, see any intermediate textbook on classical dynamics).]

So I'm a little confused at the end here... this reading is supposed to be about lumping, and yet most of the analysis is done with dimensional analysis and easy cases. The only "lumping" I remember was ignoring the implications of quantum mechanics on electrons, which seems like a necessary reduction given the complexity of quantum mechanics. Would this example be better in another section?

I was also confused by our use of "lumping" in these calculations.

I think the idea is that lumping is just one more tool that we can use.

Based on all of the examples seen here and in class, I believe that there are not many (interesting) problems that can be solved by lumping alone. But, when used in conjunction with our other tools, it complements them quite nicely.

It would have been since to have a final summary (digram!) of the different methods we used. And also the discussion of our final answer is ???

Yea I was confused as to where we used lumping in these examples. I think that lumping is used to help make the problem easier to solve (like with easy cases), but you still have to use other methods like dimensional analysis or setting the energies equal to each other in order to find a solution to the problem

I think the main idea is that for complex problems we can use lumping to help find a place to start. in this case, lumping didn't solve the problem for us but it put us in the right direction. we used the model of the potential well and then we assumed that electrons are stationary with distance r from the origin. we used this new info to come up with that graph that was missing its middle section, then we used dimensional analysis to actually come to a conclusion

Lumping as a 'tool' does seem really rather vague to me. It isn't really definable, and it seems like any way of making a problem easier can be considered 'lumping'.

I agree. This section seemed like it used techniques which we are familiar with from the class, but I'm not sure what I got out of it about lumping specifically.

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

why linearly?

I'm having trouble seeing what this has to do with lumping...

A thought on the reading: I think this reading would be really helped with a few more pictures/diagrams. Boundary layers are something that most readers do not have intuition on.

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

Read this section for Wednesay (memo due Wed at 9am) – the endpoint is golf-ball dimples.

Could you mention what boundary layers are, for the uninformed?

The boundary layer is the layer of fluid right next to the surface of whatever object you're considering. If I remember correctly, due to friction, this boundary layer can't be treated the same as the rest of the fluid, since you sort of treat that layer as though it travels with the object. (I think there is a no slip condition at the boundary)

I'm sure there is a no-slip condition.

Basically, drag to the utmost!

I'm one of those uninformed and think it'd be a good idea to clarify this too

Have patience. The section gets to it. It's easier to explain once the problem has been discussed and a diagram drawn with some context.

what paradox?

Yeah....? I thought we had successfully solved the drag problem several times so far this semester..

I agree. The paradox is explained in the next few sentences, however I was unaware that this paradox existed when I began reading the section, and it threw me a little.

Maybe combine this sentence and the one after, to read "...the drag paradox arising from analyzing drag at high Reynolds numbers..."

do you mean that it is another way to look at drag?

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

I feel like we've spent the great majority of this class just talking about drag. I thought this class was supposed to discuss several engineering disciplines... we talked about UNIX for a few classes, and then we've spent almost a month on just the concept of drag. I would appreciate it if we could move on to some other topics.

I do agree that we've spent a lot of time on drag, but at the end of the day the drag is always just the example for what we're actually being taught (dimensional analysis, lumping, etc.) I think it's nice to see the same example done using many different approximation methods so we can see exactly what's different about each of them. Maybe it just turns out that drag is one of few examples which all the different approximating methods really work for.

I do agree that we've spent a lot of time on drag, but at the end of the day the drag is always just the example for what we're actually being taught (dimensional analysis, lumping, etc.) I think it's nice to see the same example done using many different approximation methods so we can see exactly what's different about each of them. Maybe it just turns out that drag is one of few examples which all the different approximating methods really work for.

I do agree that we've spent a lot of time on drag, but at the end of the day the drag is always just the example for what we're actually being taught (dimensional analysis, lumping, etc.) I think it's nice to see the same example done using many different approximation methods so we can see exactly what's different about each of them. Maybe it just turns out that drag is one of few examples which all the different approximating methods really work for.

I do agree that we've spent a lot of time on drag, but at the end of the day the drag is always just the example for what we're actually being taught (dimensional analysis, lumping, etc.) I think it's nice to see the same example done using many different approximation methods so we can see exactly what's different about each of them. Maybe it just turns out that drag is one of few examples which all the different approximating methods really work for.

Yup, but as mentioned, drag is just one example, which we can easily use to apply a wide variety of skills and techniques in thisi class.

Btw, nice triple post there.

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

True, but some people don't care too much to be staring at discussions on drag for two months. We could hit a bunch of other cool topics and explore each with a new idea. I suppose the one positive to using drag for all of them is you see all of the approaches to one central answer, but I'm personally a fan of breadth for this class (as it seems that the idea of what we are doing is the ability to solve a large range of problems easily)

I know Reynolds number has been defined several times before, so you shouldn't define it again. But, as a non-course 2 major, I don't know what this number means or what it represents, and every time it's brought up, I have to go back through the old readings to see what it is.

It's defined as the inertial forces / the viscous forces in the fluid. It helps you understand what is physically happening in the system. High Re means that the flow is turbulent, and often little vortices are formed as a result of drag around an object.

basically think: at high reynolds number you have motion similar to a bullet flying in outer space, low reynolds number is like a bullet going through super thick molasses

rho*v*D/mu, or v*D/nu. It's a ratio of inertial to viscous forces.

It might be good here to include a table of common Reynolds numbers, so we can see what makes a high number and what it means to be a high number.

told

You've addressed this before, but I really hate the typesetting in this equation.

I think bolding the viscosity is a quick solution that would get rid of this problem

Its better to bold the velocity than viscosity since velocity is a vector (Frequently vectors will be typeset in bold)

I agree, this is still confusing without looking back.

bold one of them to distinguish the variables

In another section I read, a capital V was used for viscosity. Perhaps do that and bold it...

can you change it to rho*v*D/mu instead of vD/nu? it's often written that way instead.

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

I think this would be helpful to us if you wrote out the relation between drag coefficient and reynolds number in the equation that you have done above.

I don't really understand this comment, the equation above gives that relationship exactly.

but it doesn't explicitly relate drag coefficient and Reynolds number. It just tells you what they are in terns of variables in the equation above.

Could you explain the equation a little more? If the left side is the drag coefficient, and rv/v is the Reynolds number... Are you saying that the drag coefficient is the function of the Reynolds number? I guess I don't understand how it can be constant and a function?

So he is saying that the drag coefficient is only a function of the reynolds number; but at high reynolds number (easy to do with the viscosity goes to 0), and so if the viscosity is =0, then the Re=0, and so, the drag coefficient is now not a function of anything.

I still don't understand how we are able to take any variable from a complicated multi-variable equation, and just say that it is zero, when in reality it never is, and expect to get reasonable results from this simplified equation.

This doesn't seem to be adequately justified – I think it made more sense when we covered the topic initially.

what do you mean by as does the Reynolds number? does it vanish too? I am confused.

I think he meant that as viscosity goes to zero, some of the terms in the N-S equation vanish?

I agree that this is a bit confusing as worded. How exactly does that mean that it vanishes? Some terms vanishing means it vanishes? I suspect this actually means something about its influence, that it no longer is variable or something, but I'm not quite sure.

How does dividing by zero make something disappear? it seems to me it invalidates the relation more than anything else.

I think this is because infinity doesn't depend on v.

meaning that the coefficient of drag is constant at high Re numbers?

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

Doesn't that make it a function of infinity? f(infinity)?

I think he's saying that as the denominator goes to 0 you get f(r*nu/0)=infinity.

I agree, but that should be explained more specifically. Those of us who have no prior knowledge of these concepts have little intuition for these extreme cases.

Does this happen because the viscosity is so low that there is, in a manner of speaking, nothing to really impose drag? I'm having trouble grasping the physical implications of this particular "paradox"

yeah, someone please explain. Doesn't removing drag make this calculations useless?

I'm quite confused by this as well. here we're taking the limit as viscosity approaches zero. this makes the reynolds number very high, but how does that remove viscosity and the reynolds number from the picture?

does it imply that at low Reynolds number, drug coefficient is not constant? is it a function of viscosity then/.

Again, it'd be good here to help us get a sense of what a high Reynolds number means.

I'm pretty sure he already mentioned some examples of high and low Reynolds number fluids in a previous section. However I agree, that a quick example of something with a high Reynolds right here would be useful, especially for people like me who don't really know anything about fluid mechanics.

this should either be " at A high reynolds number" or "at high reynolds numberS"

This is what I would expect given how the Reynolds number affects the mechanics of fluids.

I agree but I'm not all that comfortable with how we got there.

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

I see you've tried to differentiate between velocity and viscosity by bolding velocity, and that seems to work fine. Do you think you could make this change to the variable sin the written paragraphs too?

I assumed the bolding was to indicate it's a vector, not to differentiate it from nu? If that's not true, than it's also a bit confusing to have it bolded.

I think seeing the notation in this equation made me more confused between nu and v. But I understand the concept.

If person 1 is correct and the bold is to differentiate, then I think it works nicely.

I can't remember earlier notation in the previous chapters but if bolding is used to identify vectors that is another issue I hadn't considered with using bold to differentiate the two symbols. Otherwise I think bolding does the job very nicely

I believe nu is written as a curvy v or a v in italics and velocity is a regular v. I think the bold is to show that it's a vector. If you look at the second term on the right of the equal sign it's an italics v multiplied by the gradient squared multiplied by a bold v.

Maybe a quick note to let the reader know you are bolding velocity to distinguish it would go a long way.

So velocity can be considered a vector (or each cartesian coordinate can be considered independently.. but the grads would change to partial derivatives), so that interpretation wouldn't be incorrect. I think it also serves to differentiate from nu.

You can also try to use the little vector arrow above the v.

using viscosity has always bugged me a little bit I feel like we would want to use velocity as the variable rather than viscosity. which is what we did in 2.005/2.006 and it seems to make for physical sense

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \quad (7.8)$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \quad (7.9)$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

The bolded v here is velocity correct? If velocity is bolded here than it may help to bold velocity above in the Reynolds number equation to help solidify which variables are what

The bolded v here does stand for velocity. In the Reynolds equation, the v, which also stands for velocity is not bolded, and I agree that it might be slightly confusing, but I think the bolded variables are vectors. Perhaps there should be a note about this in the text since so many people are confused about this.

this sentence syntax makes no sense to me.

Yeah – it uses two different appositive constructions (commas and dashes), within a parenthetical statement. Probably too much to be going on in one sentence.

yea the phrase makes sense to me without the – the pressure – and once it's added, i'm not sure what it's referring to

Will this book have an appendix for the requisite mathematical terms and concepts, or will it not be self-contained?

I think terms like gradient and scalar are pretty basic, and information about them is readily available in most math text books of the internet. There are definitely more complicated terms and topics discussed that I would love to see external links to where I can find more information as I often mind myself looking into the things discussed in these chapters more on my own time

Why can't it dissipate energy?

It can't dissipate energy because it is a conservative force, which means essentially that it is reversible without loss of energy. Think of it like gravity: if you take a ball at a height h1 and bring it to a height h2, it loses some amount of energy, but if you repeat the experiment it will always lose the same amount of energy - nothing is dissipated to the surroundings.

Based on my intuition I would expect the drag coefficient to approach a really large value so I agree the drag coefficient cannot approach a constant

oh..this make more sense!

so much better now that this was explained - had no idea where there was a paradox before.

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \tag{7.8}$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \tag{7.9}$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!* In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

wait...are you saying it should or should not be 0?

I think he's saying that we know it's not zero, but if you make that assumption this equation implies that it should be zero.

He's saying that the equation says that it should be, but we know from real life that it isn't.

I thought we just said it was infinity?

No, we were taking the Reynold's number to infinity to have a roughly constant drag coefficient.

where does the discrepancy lie? Is there another area where drag coefficient can be measured as a constant?

I think the discrepancy is that according to the equation, drag coefficient (and thus drag) should approach zero as viscosity approaches zero, but in real life drag is never 0

Interesting...I wonder how we resolve this.

I would never have figured out any of this on my own, I would have just ignored N-S!

He also pointed this out on the graph though. The line goes relatively flat, but it's not flat at zero.

It took a while to explain what you meant in the intro.

Yeah, but I am glad you did. When I read it in the introduction of the section, I had no idea what it was and thought I was about to not understand this example at all. So I am glad it is included.

I would love to know why that is a paradox

basically because the navier stokes equation (which we assume to be correct) says that if viscosity goes to zero then the drag coefficient should approach the constant zero. However that doesn't happen in real life. So it is a paradox because both of these can't be true.....

7.4 Boundary layers

Boundary layers, which are the final example of lumping, will help us explain the drag paradox. That paradox arises in analyzing drag at high Reynolds numbers (the usual case in everyday fluid motions). Dimensional analysis tells us that the drag coefficient c_d is a function only of the Reynolds number:

$$\frac{F_d}{\frac{1}{2}\rho v^2 A} = f\left(\frac{rv}{\nu}\right), \quad (7.8)$$

where $F_d/\frac{1}{2}\rho v^2 A$ is the drag coefficient, and rv/ν is the Reynolds number. To make a high Reynolds number, take the limit in which the viscosity ν approaches 0. Then viscosity vanishes from the analysis, as does the Reynolds number. The result is that the drag coefficient is a function of nothing – in other words, it is a constant. So far, so good: Empirically, at high Reynolds number, the drag coefficient is roughly constant and independent of the Reynolds number.

The paradox arises upon looking at the Navier–Stokes equations:

$$(\mathbf{v}\cdot\nabla)\mathbf{v} + \frac{\partial \mathbf{v}}{\partial t} = -\frac{1}{\rho}\nabla p + \nu\nabla^2\mathbf{v}. \quad (7.9)$$

When the viscosity ν goes to zero, the $\nu\nabla^2\mathbf{v}$ viscous-stress term also vanishes. However, the viscous-stress term is the only term that dissipates energy. (The pressure term ∇p is, like gravity, a conservative force field because it is the gradient of a scalar – the pressure – so it cannot dissipate energy.) Without the viscous-stress term, there can be no drag! So, at high Reynolds number, the drag coefficient should approach a constant *but that constant should be zero!*. In real life, however, the drag is not zero; this discrepancy is the drag paradox.

Mathematically, the paradox is a failure of two operations to commute. The two operations are (1) solving the Navier–Stokes equations and (2) taking the viscosity to zero. If we first solve the Navier–Stokes equations, then we find that the drag coefficient is nonzero and roughly constant (i.e. independent of Reynolds number). If we then take the viscosity to zero (by taking the Reynolds number to infinity), no harm is done. Because the drag coefficient is roughly independent of Reynolds number, the drag coefficient remains nonzero and roughly constant.

are there other examples of this? how do we know if it's this or just that we messed up in our analysis?

Agreed, I don't know if I would be able to recognize this paradox myself.

Also agreed. I would be interested to know if there are other such paradoxes in physics/mechanics.

I've run into this problem before for other equations, and realized that it's usually better to solve the equation before plugging in an edge case.

What other examples do you mean?

I definitely agree, but there are cases where it's fine to assume limits before solving. I wonder if there's any more systematic way of knowing when you can and when you can't?

I think this sentence is confusing?

I feel like here is where I finally start to see where boundary layers are talked about, but a definition/earlier introduction might be more helpful for directing this section

wait these seem like they work in harmony instead of dissonance – what am i missing?

Continue reading into next paragraph. The paradox arises then.

The paradox is described in the next paragraph.

wait really? can we see this out in class today? i don't quite get how now hard is done this way

How is the drag coefficient roughly independent but can be a function of Reynolds number?

After a point, it no longer varies significantly with the increasing Reynolds number. It is a function to a point, but becomes roughly independent.

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

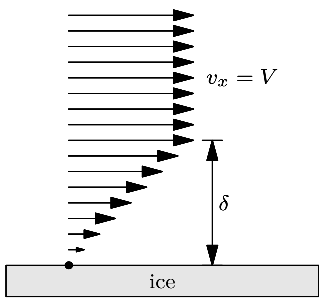

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

Comments on page 2

It's so interesting to me that you can use the exactly right equations and still come out with the wrong answer simply due to the order in which you solve them.. Do you have any hints on how to know which equation to approach first (in a general case)?

I'd say it's a decision about when to throw out information. Here we were foiled because we made our lossy assumption first (before using N-S which requires a viscosity).

In other cases, making lossy assumptions first will make problems much more convenient to solve, so it's probably a judgment call. My approach would be try making the assumption first, and sense-check the result to see if your assumption was 'legal'

I think this is a lot of text to explain math. I was having trouble following, especially after the page break.

I agree...it was a bit difficult to follow math via text. I think a small diagram or display of the two methods might help visualize the idea.

Yes it is a lot of text, but it seems necessary to explain this. as mentioned a comparison of the two methods (and how the viscous-stress term drops out) would help to tidy this up

I was totally confused until this point...I think that this would work better with a picture/graph or actually writing the equations.

I agree, seeing the two methods of solving this problem would help a lot here.

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

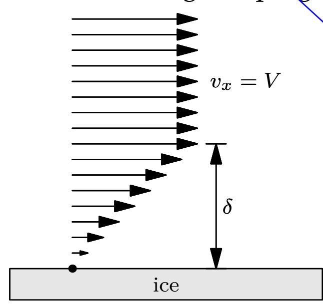

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

Even after this is explained, I still find it incredulous that such important math equations work out differently depending on order!

It has to do with the lossy approximation of viscosity going to zero and striking out that term from N-S before solving it.

I think this relates to the problems one can run into if you throw out the zeroes too early in certain problems.

Yeah, it has to do with the order of the problem (wow, English is confusing... by order of the problem I mean removing the highest-order derivative not order of solving equations). THis is explained in the next paragraph.

I'd go with 'shocking' over 'incredulous' ... I understand why, it's just that it's a little hard to believe without really thinking about it.

yeah i think this has to do with our sketchy math at the beginning more than anything else.

are there other equations like this?

yeah I would be really curious to hear about some other example of equations like this in which the solution depends on the order!

Just to clarify, you follow this with an explanation but it's not immediately clear if that's a definition of what a singular limit is here, or just this specific case with drag...

Perhaps change the next sentence to "It is the limit that..." would make it more clear.

I agree with the previous statement; I'm not sure if it's a definition for singular limit, but I have to assume it is because I don't know what a singular limit is (though I think it removes the highest-order-derivative...)

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

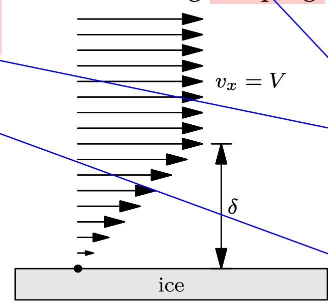

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

I thought this idea was really cool - that the operations can't commute because you're reducing the dimensions of the equation by switching the orders.

That does seem really crazy- by doing the math in a certain order you can actually obtain a contradiction.

Does this relate to the problem with taking a square root in an equation? For instance, reducing xˆ2 = 4 to x=2. I understand its not an issue of limits, but its a mathematical operation that is communicable.

That's really cool. The first thing I thought of was operators in quantum mechanics, but I never really got that some operators can't commute because the commutation changes the order of the derivative until I read this.

This sentence is confusing as it is worded now... maybe say, "this qualitative change in the order of the equation produces a qualitative change in behavior."

I have a feeling its a type. Maybe he meant to say that "this quantitative change produces a qualitative..."

Ahh, so this is the origin of the boundary layer. Very cool!

Is this an example of a mathematician using physical reasoning to solve an equation?! Gasp!

good indicator of next steps.

how true is the no slip condition?

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

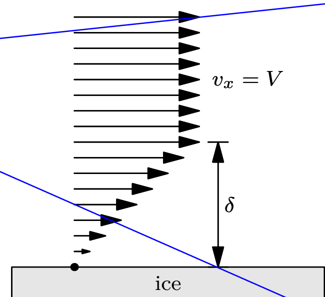

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

Why is the type of surface not a significant factor in the size/significance of the boundary layer. I would think that more slippery objects would cause the boundary layer to be much thinner, and therefore change the drag significantly. Is this the case or is it not really that important?

"Slippery object" isn't really an appropriate concept. First of all, we are assuming what's called a "no-slip" condition. This means that the fluid that is in contact with the object is not moving with respect to the object. Slightly more explanation here: http://en.wikipedia.org/wiki slip_condition

Second, given that no-slip condition, any non-geometric properties of the object won't have an effect. Fluid that is moving past other fluid is only affected by the fluid's own viscosity, not by any "slipperiness" of the object. I kind of like this example: http://en.wikipedia.org/wiki/Couette_flow

Maybe you could say "In a fluid there is viscosity, and the fluid is..." instead of the current phrasing which requires an explanation in parenthesis.

I don't think that change would make it better...it seems a bit awkward. Why not just say "All fluids have viscosity and is at rest next to a solid object."

OK, I was just trying to make a helpful suggestion. Since you want to get picky about it, your sentence isn't even grammatically correct. The sentence: "All fluids have viscosity and ARE at rest..." has the appropriate plurality.

And I am just giving my views on it. There's no needs to start an argument.

hahaha. fight on nb. but seriously...you go back and read your replies?? weird.

didnt we just make a whole argument assuming that the viscosity didn't exist??? i dont really understand why we didnt start with this statement.

I think the point is to show how the previous model does not explain the paradox - that we need to use another method to figure this out.

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

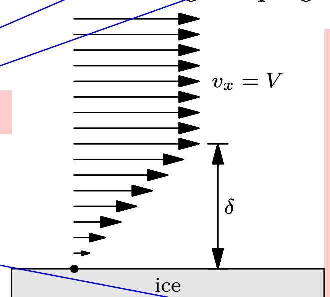

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

Is the no-slip condition true of all fluids in the physical world, or is at an analytical shortcut - like "imagine an object falling in a vacuum"?

I think this has to be a true condition, or at least a good approximation to reality. It seems to me that molecular interactions between the "object" and the "fluid" would cause the object to drag a very thin layer of the fluid along with it, and this layer would then have zero velocity with respect to the object.

This scenario seems like it has a large number of application to everyday events.

This is a great example because it's so easy to picture and everyone is able to relate it

This is a great description! I can definitely visualize the wind and the frozen lake interaction.

Agreed. I feel this way about a lot of explanations in this class - I wish they'd be given earlier

excellent diagram- why wait till now to show it? maybe have a diagram of the paradox. visuals i feel are key in explaining this content

Yeah I really like this picture. It's really helpful in understanding this example. It reminds me of the difference between Newtonian and relativistic mechanics.

Agreed. I think its a great diagram and these comments just show that maybe more diagrams are needed in areas that involve heavy mathematical explanations.

I'd say: "...frozen lake. Despite the wind, the air has zero velocity in the horizontal direction just above the ice, and . (In the vertical..."

I had to re-read this paragraph to remember that you were taking about just above the ice...switching the order would have really helped.

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

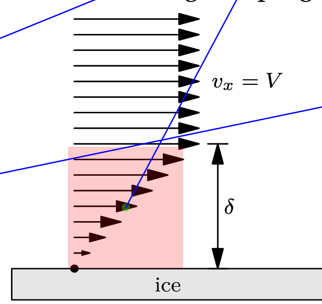

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

Is this diagram showing the linear function that we estimated with lumping or is this a more accurate function?

This is the lumping approximation. And the consant part above is also an example of lumping.

This diagram is excellent for understanding our lumping model.

I agree that this is an excellent visual of our lumping model, but it may help to put the equation for the boundary layer region the same way the our upper part is v_x = V

I like the picture here, it helps make these sentences very easy to understand.

Definitely. I was somewhat picturing a similar graph in my mind, but that one here is quite helpful in understanding the phenomenon.

I agree, this is similar to pictures we used in 2.005 and pictures like this probably could have been used in earlier sections as well.

I don't understand this... why does the air have zero horizontal velocity next to the ice?

This was new information for me too. Is this something we could have observed on our own? How big is the boundary layer in this case?

why is the speed slower (or zero) closer to the ice? is it just because the ice is cold that it slows down the wind somehow?..or..what.

I am confused by this as well. conceptually this may need a little more explanation. how can the horizontal velocity be zero by the ice, and what causes the boundary layer?

I think I left out an important intuition that belongs the explanation, namely that viscosity prevents the horizontal flow speed from changing too rapidly. So, the flow speed is zero next to the ice and is at full speed far away, and it varies smoothly in between thanks to viscosity. I'll explain that more in lecture...

This is really cool. Another interesting fact that I never knew.

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

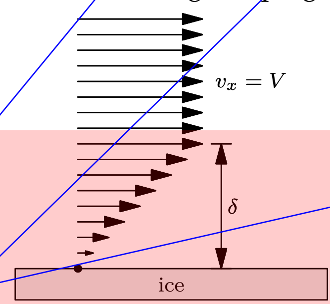

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

I love this example- I was starting to get lost, but I understand exactly what you mean now with this application

I agree. This simple example, coupled with the diagram on the right, help clarify some confusing points above.

I also agree. It's easy to get lost in the math, but the example here and those in class make it all click.

I like this explanation in the context of the example, but it may be helpful to have a more formal explanation of what a boundary layer is as well.

I don't think it should be too formal. If you read on, he kind of supplements his explanation here with more examples and explanation. I think it's sufficient.

This is I think as formal as you can get by example...I like it

I feel like the way he explains it is great because it helps you remember it in a visual way. I personally probably would skim over a formal definition and forget it, but I won't forget this.

I thought this explanation was as good as it could be.

i understand the concept of the boundary layer but i don't understand how this solves the paradox

"Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity)."

Why not?

yeah this seems like a strange digression.

Maybe because of interference from the vertical velocity? Just a wild guess.

This is really where this section start making sense to me, the example is very clear, and I finally understand was is meant by boundary layer

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

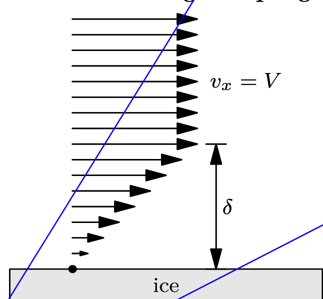

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

Do you mean just in the boundary layer it doesn't reach V, or is there some bigger point I'm missing?

I think he means in the boundary layer, otherwise V wouldn't exist at all.

It looks like just in the boundary layer in his diagram

I agree there is some slight confusion. But I do think he means just in the boundary layer.

I thought this had something to do with the air resistance.

It's analogous to an object never reaching terminal velocity – it just approaches that speed more and more as it falls. Similarly, as you go farther away from the ice, the speed gets closer to V but it never reaches it fully. The boundary layer is, intuitively, the region over which it mostly gets to V.

I'm not sure this sentence to the left is actually a sentence. I'm not sure if it was meant this way, or it is meant to be one.

I know we're assuming the relationship is linear for the sake of a lumping approximation, but what shape does this relationship take empirically?

It's been about a year since I did boundary layers, but I'm pretty sure it resembles the shape of the f(x) = 1 - (e^-x) function.

That sounds right. It's also commonly approximated by a parabola instead of a line. The parabola gives a quite good fit to the actual velocity profile, without adding too much complexity as to make the governing equations unsolvable.

This is a nice refresher on material from 2.005

So in general a boundary layer is the space that doesn't reflect the surrounding velocity of air?

I say: "...viscosity. The dimensions of viscosity (L2T-1), along with a bit of dimensional analysis, will help us estimate the thickness. To make..."

where is the dimensional analysis for this? I am a little confused trying to follow what is going on

Now imagine applying the two operations in the reverse order. Taking the viscosity to zero removes the viscous-stress term from the Navier–Stokes equations. Because that term is the only loss term, solving these simplified equations – called the Euler equations – produces a solution with zero drag. The mathematical formulation of the drag paradox is that the solution depends on the order in which one applies the operations of solving and taking a limit.

This paradox, which remained a mystery for over 100 years, was resolved by Prandtl in the early 1900s. The resolution identified the fundamental problem: that taking the viscosity to zero is a *singular limit*. That limit removes the highest-order-derivative from the differential equation, so it changes the equation from a second- to a first-order equation. This qualitative change produces a qualitative change in behavior: from nonzero drag to zero drag. To handle this change properly, Prandtl devised the concept of a boundary layer – which we can understand using lumping.

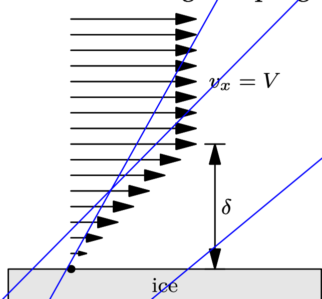

The explanation begins with the no-slip condition: In a fluid with viscosity (which means all fluids), the fluid is at rest next to a solid object. As an example, imagine wind blowing over a frozen lake. Just above the ice, and despite the wind, the air has zero velocity in the horizontal direction. (In the vertical direction, the velocity is also zero because no air enters the ice – that requirement is independent of the no-slip condition.) Far above the ice, the air has the speed of the wind. The boundary layer is the region above the ice over which the horizontal velocity v_x changes from zero to the wind velocity. Actually, the horizontal velocity never fully reaches the full wind velocity V (called the free-stream velocity). But make the following lumping approximation: that v_x increases linearly from 0 to V over a length δ – the thickness of the boundary layer.

The boundary layer is a result of viscosity. The dimensions of viscosity, along with a bit of dimensional analysis, will help us estimate the thickness δ. The dimensions of ν are L^2T^{-1}. To make a thickness, multiply by a time and take the square root. But where does the time come from? It is the time which the fluid has been flowing over the object. For example, for a golf ball moving at speed v, the time is roughly $t \sim r/v$, where r is the radius of the ball. Then

Why is the L dimension squared?

I think that's the viscosity v, not velocity.

In case you knew it was viscosity, it makes sense because the area matters when two things rub against each other.

This feels a bit arbitrary.

Now, given that v=velocity is an important parameter, why not just divide by velocity? [nu/v]=L. Why would we guess to include a diameter as well? Or at least, why would we ignore the dimensionless group v*delta/nu?

Im a little confused about the time dimension here.

It is very unclear to me what this time is. Does this only apply to the fluid at the bottom of the boundary layer, or is it the average time that the average fluid involved in the boundary layer takes to get around the object, or something else entirely. I would appreciate a little more explanation about this.

Also, if we were to take the frozen lake example, do we use a T from the time the ice was formed to the time of measurement? But I guess this might make sense, for if you take antarica, a larger time, will decrease the boundary layer, making surface winds very high, which is true...

what time are looking at? what is the reasoning behind choosing r?

Nevermind, I understand now that you want the time it takes to pass over the object.

$$\delta \sim \sqrt{\nu t} \sim \sqrt{\frac{\nu r}{v}}. \qquad (7.10)$$

Relative to the size of the object r, the dimensionless boundary-layer thickness δ/r is

$$\frac{\delta}{r} \sim \sqrt{\frac{\nu}{rv}}. \qquad (7.11)$$

The fraction inside the square root looks familiar: It is the reciprocal of the Reynolds number Re! Therefore

$$\frac{\delta}{r} \sim \mathrm{Re}^{-1/2}. \qquad (7.12)$$

For most everyday flows, $\mathrm{Re} \gg 1$, so $\mathrm{Re}^{-1/2} \ll 1$. The result is that the boundary layer is a thin layer.

This thin layer will resolve the drag paradox. The intuition is that the boundary layer separates the flow into two regimes: inside and outside the boundary layer. Inside the boundary layer, viscosity has a large effect on the flow. Outside the boundary layer, the flow behaves as if viscosity were zero; there, the flow is described by the Euler equations (by the Navier–Stokes equations without the viscous-stress term). However, the boundary layer does not stick to the object everywhere. Generally, it detaches somewhere on the back of the object. Once the boundary layer detaches, a wake – the region behind the detached layer – is created, and the wake is turbulent. The wake has high-speed and therefore low-pressure flow: Bernoulli's principle says that pressure p and velocity v are related by $p + \rho v^2/2 = \text{constant}$, so high v implies low p. Therefore, the front of the object experiences high pressure and the back experiences low pressure. The result is drag. Intuitively, the drag coefficient is the fraction of the cross-sectional area covered by the turbulent wake.

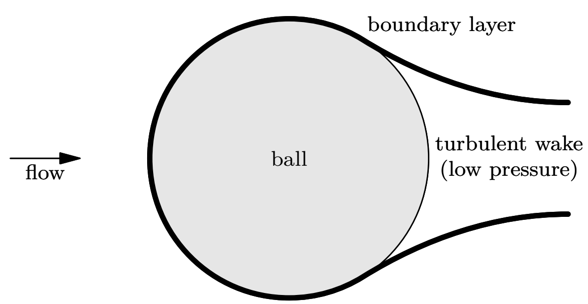

$$\delta \sim \sqrt{\nu t} \sim \sqrt{\frac{\nu r}{v}}. \quad (7.10)$$

Relative to the size of the object r, the dimensionless boundary-layer thickness δ/r is

$$\frac{\delta}{r} \sim \sqrt{\frac{\nu}{rv}}. \quad (7.11)$$

The fraction inside the square root looks familiar: It is the reciprocal of the Reynolds number Re! Therefore

$$\frac{\delta}{r} \sim \mathrm{Re}^{-1/2}. \quad (7.12)$$

For most everyday flows, $\mathrm{Re} \gg 1$, so $\mathrm{Re}^{-1/2} \ll 1$. The result is that the boundary layer is a thin layer.

This thin layer will resolve the drag paradox. The intuition is that the boundary layer separates the flow into two regimes: inside and outside the boundary layer. Inside the boundary layer, viscosity has a large effect on the flow. Outside the boundary layer, the flow behaves as if viscosity were zero; there, the flow is described by the Euler equations (by the Navier–Stokes equations without the viscous-stress term). However, the boundary layer does not stick to the object everywhere. Generally, it detaches somewhere on the back of the object. Once the boundary layer detaches, a wake – the region behind the detached layer – is created, and the wake is turbulent. The wake has high-speed and therefore low-pressure flow: Bernoulli's principle says that pressure p and velocity v are related by $p + \rho v^2/2 = \text{constant}$, so high v implies low p. Therefore, the front of the object experiences high pressure and the back experiences low pressure. The result is drag. Intuitively, the drag coefficient is the fraction of the cross-sectional area covered by the turbulent wake.

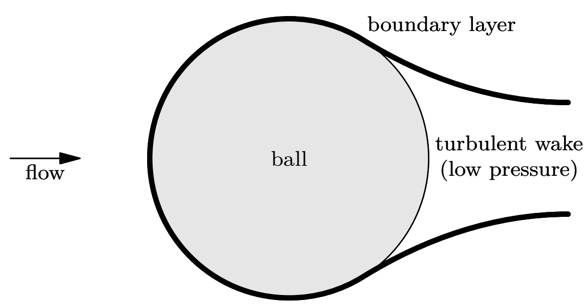

haha, absolutely amazing.

I never would have guessed that the Reynolds number would be so important and useful in all of these applications.

Also, I'm still impressed that we continue to determine all of these things via dimensional analysis.

Totally agree, the dimensional analysis techniques we've learned in this class have helped me across the board in all of my other classes.

I suppose that depends on what major you are. I haven't been able to apply any of this to my other classes. I wish I could...

What major are you that none of this is applicable?

I agree with the earlier comment on not being able to apply any of this material.... All this stuff we've been learning is very interesting, but as a course 6 major, it simply has nothing to do with our curriculum.

I find that hard to believe. Surely, as a course 6 major, you've had to write unit tests before? Divide and conquer, easy cases, and lumping are applicable there.

Perhaps in your future career, as a course 6 major, you will be asked by your client/employer to develop a system capable of handling hundreds or thousands of users. Divide and conquer, along with some lumping, could allow you make a rough estimate as to how scalable your system needs to be.

Perhaps you are right in that you may never have to work with a Reynolds number in your curriculum or career, but it's a way of thinking, not knowledge of this specific example, that this class is trying to develop. :)

As well as all the stuff we did at the beginning with unix. Or the stuff we have been doing with circuitry.

$$\delta \sim \sqrt{\nu t} \sim \sqrt{\frac{\nu r}{v}}. \qquad (7.10)$$

Relative to the size of the object r, the dimensionless boundary-layer thickness δ/r is

$$\frac{\delta}{r} \sim \sqrt{\frac{\nu}{rv}}. \qquad (7.11)$$

The fraction inside the square root looks familiar: It is the reciprocal of the Reynolds number Re! Therefore

$$\frac{\delta}{r} \sim \mathrm{Re}^{-1/2}. \qquad (7.12)$$

For most everyday flows, $\mathrm{Re} \gg 1$, so $\mathrm{Re}^{-1/2} \ll 1$. The result is that the boundary layer is a thin layer.

This thin layer will resolve the drag paradox. The intuition is that the boundary layer separates the flow into two regimes: inside and outside the boundary layer. Inside the boundary layer, viscosity has a large effect on the flow. Outside the boundary layer, the flow behaves as if viscosity were zero; there, the flow is described by the Euler equations (by the Navier–Stokes equations without the viscous-stress term). However, the boundary layer does not stick to the object everywhere. Generally, it detaches somewhere on the back of the object. Once the boundary layer detaches, a wake – the region behind the detached layer – is created, and the wake is turbulent. The wake has high-speed and therefore low-pressure flow: Bernoulli's principle says that pressure p and velocity v are related by $p + \rho v^2/2 = \text{constant}$, so high v implies low p. Therefore, the front of the object experiences high pressure and the back experiences low pressure. The result is drag. Intuitively, the drag coefficient is the fraction of the cross-sectional area covered by the turbulent wake.

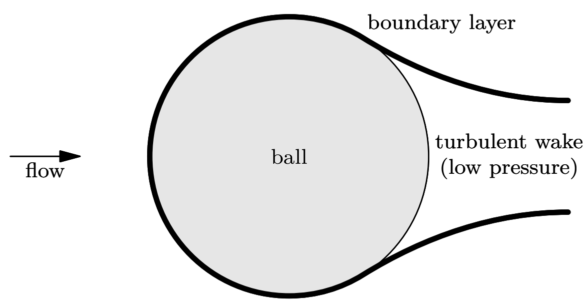

Yeah, I understand this class is trying to develop a "way of thinking"... my point is that, as a course 6 person, this class hasn't taught me anything I didn't already know. If you can't understand a basic concept like D&C in course 6, you're screwed. But that's why I learned it and UNIX and circuitry 3 years ago. Most of the "handy tricks" that he's shown us that aren't taught in other classes are for course 2 majors. So yeah, nothing I've LEARNED in this class has helped me. Because discussing the same drag problem for a month doesn't help course 6 people.

This is awesome. Great movement from the earlier sections to here; and then right into the golf balls. Good section.

What kinds of topics and problems are you doing in the course 6 courses? (And, are you 6-1, 6-2, or 6-3?)

I might like to see this particular problem with the golf ball solved so we can see what a reasonable number is... I know we already know delta and r and could do this ourselves but it would be fast and cool to see the calculation done here to show that this statement is valid.

I intuitively would hope this to be the case.

I would suspect that this would be true with air, but I would be curious to see how large the boundary layer is with objects flowing in water, a very everyday flow. I would guess that it might be quite large.

I would love to see the thickness of this layer calculated, and how it is affected by different changes in viscosity or speed

Oh, that's a good point. Those would be interesting calculations to see for sure.

i'd say "the boundary layer is thin"

..."the boundary layer is a thin layer" sounds funny

how does this resolve the paradox?

$$\delta \sim \sqrt{\nu t} \sim \sqrt{\frac{\nu r}{v}}. \qquad (7.10)$$

Relative to the size of the object r, the dimensionless boundary-layer thickness δ/r is

$$\frac{\delta}{r} \sim \sqrt{\frac{\nu}{rv}}. \qquad (7.11)$$

The fraction inside the square root looks familiar: It is the reciprocal of the Reynolds number Re! Therefore

$$\frac{\delta}{r} \sim \mathrm{Re}^{-1/2}. \qquad (7.12)$$

For most everyday flows, $\mathrm{Re} \gg 1$, so $\mathrm{Re}^{-1/2} \ll 1$. The result is that the boundary layer is a thin layer.

This thin layer will resolve the drag paradox. The intuition is that the boundary layer separates the flow into two regimes: inside and outside the boundary layer. Inside the boundary layer, viscosity has a large effect on the flow. Outside the boundary layer, the flow behaves as if viscosity were zero; there, the flow is described by the Euler equations (by the Navier–Stokes equations without the viscous-stress term). However, the boundary layer does not stick to the object everywhere. Generally, it detaches somewhere on the back of the object. Once the boundary layer detaches, a wake – the region behind the detached layer – is created, and the wake is turbulent. The wake has high-speed and therefore low-pressure flow: Bernoulli's principle says that pressure p and velocity v are related by $p + \rho v^2/2 = \text{constant}$, so high v implies low p. Therefore, the front of the object experiences high pressure and the back experiences low pressure. The result is drag. Intuitively, the drag coefficient is the fraction of the cross-sectional area covered by the turbulent wake.

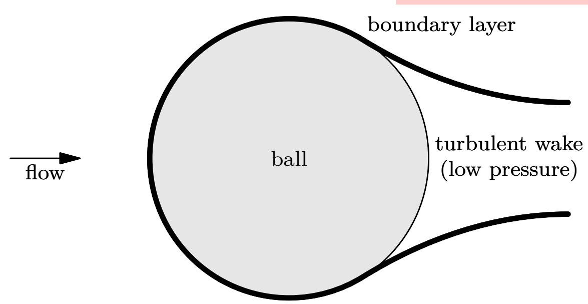

Does the surface friction of the object also play a role here, or is the relationship always no-slip?

It's always a no-slip requirement regardless of material or finish.

I don't really understand this. can you clarify what you mean by 'sticking', and why this happens?

I sort of understand the idea for the ball, but what about for the ice example? Where is the wake?

I agree it's pretty easy to see for the ball, but what about in other situations?

really helpful to have the picture of the flow around the ball, otherwise I might have been more confused.

This is sort of a deceptive use of the term "turbulent." Technically any flow with Re > 2300 is turbulent (not laminar). We're assuming that the flow outside the boundary layer is turbulent, but I think you mean to distinguish the flow in the wake as having vortices.

This is very interesting, and another great intuitive explanation for drag. I like this idea of taking a concept and studying it from a variety of different perspectives throughout the term.

The fraction explanation of the drag coefficient is also very easy to visualize.

I agree.

I enjoy seeing the same concept explained in a variety of similar, yet different, ways. It really helps hammer out the uncertainties and increases understanding.

We just derived Bernoulli's equation in 2.005, and it's great to see it broken down and applied in a different way. Really helps to get a better understanding of the concept

That certainly wasn't my intuition...because I still don't have any about drag. Can you explain this more?

I's really like an extra sentence or 2 to explain this. I think it makes sense in that something with a really turbulent wake also have a really high drag force.

I agree, just thinking about the situation and the idea of a turbulent wake, this makes it all easier to understand.

$$\delta \sim \sqrt{\nu t} \sim \sqrt{\frac{\nu r}{v}}. \quad (7.10)$$

Relative to the size of the object r, the dimensionless boundary-layer thickness δ/r is

$$\frac{\delta}{r} \sim \sqrt{\frac{\nu}{rv}}. \quad (7.11)$$

The fraction inside the square root looks familiar: It is the reciprocal of the Reynolds number Re! Therefore

$$\frac{\delta}{r} \sim \mathrm{Re}^{-1/2}. \quad (7.12)$$

For most everyday flows, $\mathrm{Re} \gg 1$, so $\mathrm{Re}^{-1/2} \ll 1$. The result is that the boundary layer is a thin layer.

This thin layer will resolve the drag paradox. The intuition is that the boundary layer separates the flow into two regimes: inside and outside the boundary layer. Inside the boundary layer, viscosity has a large effect on the flow. Outside the boundary layer, the flow behaves as if viscosity were zero; there, the flow is described by the Euler equations (by the Navier–Stokes equations without the viscous-stress term). However, the boundary layer does not stick to the object everywhere. Generally, it detaches somewhere on the back of the object. Once the boundary layer detaches, a wake – the region behind the detached layer – is created, and the wake is turbulent. The wake has high-speed and therefore low-pressure flow: Bernoulli's principle says that pressure p and velocity v are related by $p + \rho v^2/2 = \text{constant}$, so high v implies low p. Therefore, the front of the object experiences high pressure and the back experiences low pressure. The result is drag. Intuitively, the drag coefficient is the fraction of the cross-sectional area covered by the turbulent wake.

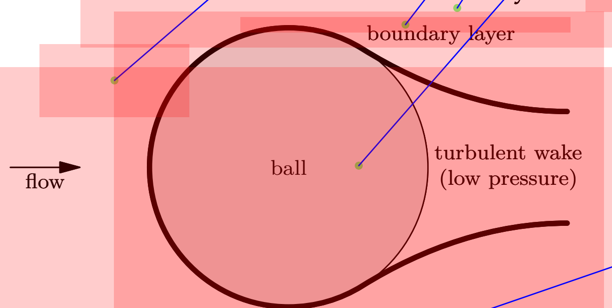

and turbulent is related to a higher reynolds number?...how does that relate to this sentence?

I'm not 100% sure what covering a cross-sectional area with turbulent wake means. It would make more sense to me for it to be the fraction of the _perimeter_ covered by turbulent wake.

how would dimples on a golf ball affect the size of the boundary layer on a golf ball?

to make the diagram perfectly clear, it'd be helpful to have the boundary layer specified as the thick line (arrow + circle or something)

I really like this image. I'm not sure its necessary, but its very clean and dovetails nicely with the previous paragraph

I think it fits well here because it helps visualize where the boundary layer would detach on an object, and how this causes the difference in pressures that are responsible for drag

I think the picture is necessary. Overall, this section was a VERY technical read, and having some images and diagrams definitely appeals to different types of learners. Having the picture here doesn't hurt anyone so it should definitely be kept in.

Yeah, I personally really like this as well. It's definitely clean as you described.

I like the picture, however, could we have a real life example of this? Maybe we could talk about it or see it in class.

This reminds me of the golf ball question from the diagnostic... do the dimples increase boundary layer thickness making the drag lower?

So after we did our diagnostic, I decided to read into it, and found that there actually is a wikipedia article on golf balls:

http://en.wikipedia.org/wiki/Golf_ball

So it seems like golf balls are made with dimples to induce a boundary layer and reduce the drag force.

Helpful diagram.

Looks great and gets the point across perfectly. If it wasnt here I probably would've been trying to draw one in the margins.

That interpretation of the drag coefficient leads to an explanation of why the drag coefficient remains roughly constant as the Reynolds number goes to infinity – in other words, as the flow speed increases or the viscosity decreases. In that limit, the boundary-layer detachment point shifts as far forward as it goes, namely to the widest portion of the object. Then the drag coefficient is roughly 1.

This explanation mostly accounts for the high-Re data on drag coefficient versus Reynolds number. Here is the log–log plot from Section 5.6:

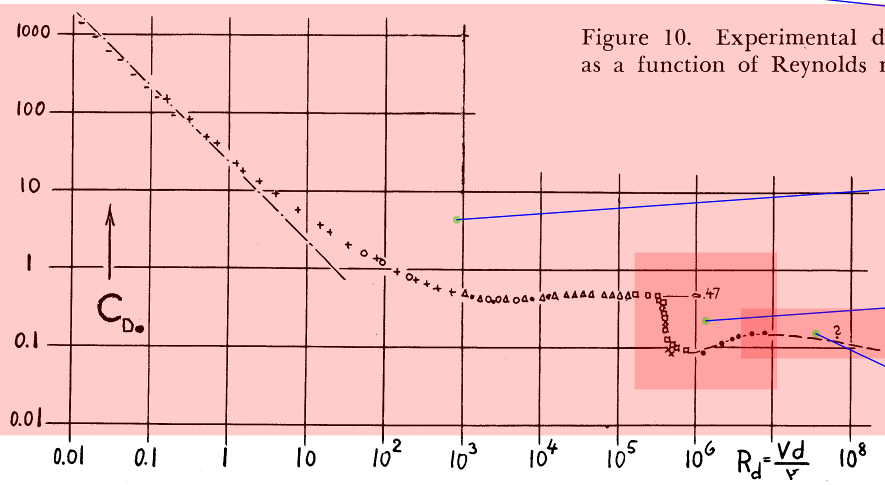

Figure 10. Experimental d
as a function of Reynolds ı

The drag coefficient is roughly constant in the Reynolds-number range 10^3 to (almost) 10^6. But why does the drag coefficient drop sharply around $\text{Re} \sim 10^6$? The boundary-layer picture can also help us understand this behavior. To do so, first compute Re_δ, the Reynolds number of the flow in the boundary layer. The Reynolds number is defined as

$$\text{Re} = \frac{\text{typical flow speed} \times \text{distance over which the flow speed varies}}{\text{kinematic viscosity}}. \quad (7.13)$$

In the boundary layer, the flow speed varies from 0 to v, so it is comparable to v. The speed changes over the boundary-layer thickness δ. So

$$\text{Re}_\delta \sim \frac{v\delta}{\nu}. \quad (7.14)$$

Because $\delta \sim r \times \text{Re}^{-1/2}$,

Comments on page 4

where is the lumping?

This makes sense for a sphere. I'm not 100% sure about this being true for more complicated shapes.

I like that you show the plot again here instead of making the reader flip back and find it.

Agreed, definitely convenient for an ebook

What is the rationale for a log-log plot?

I guess this part answers my question earlier.

Why is the graph of Reynolds numbers versus distance supposed to have constant slope under the log-log scale?

So I've already been taught about what happens here physically in previous classes, but every time I see this graph I'm still perplexed by this little dip. I know why it does it, but it just doesn't seem "right" that this is just what happens naturally.

What did your previous classes say about this?

How do we explain this behavior?

I might be wrong, but it looks like it becomes constant again like it did at 10^3 to 10^6, except no fall again at an even higher Reynolds number. Anyone else have thoughts?

I'd still like to see better axis labels here...put them actually outside of the graph.

I find this very frustrating. If the graph is actually a log-log plot then the high end is approximately 2*10^5 ... much closer to 10^5 than 10^6

that often bugs me too, but "closer" could be used here in a geometric (as opposed to arithmetic) sense.

That interpretation of the drag coefficient leads to an explanation of why the drag coefficient remains roughly constant as the Reynolds number goes to infinity – in other words, as the flow speed increases or the viscosity decreases. In that limit, the boundary-layer detachment point shifts as far forward as it goes, namely to the widest portion of the object. Then the drag coefficient is roughly 1.

This explanation mostly accounts for the high-Re data on drag coefficient versus Reynolds number. Here is the log–log plot from Section 5.6:

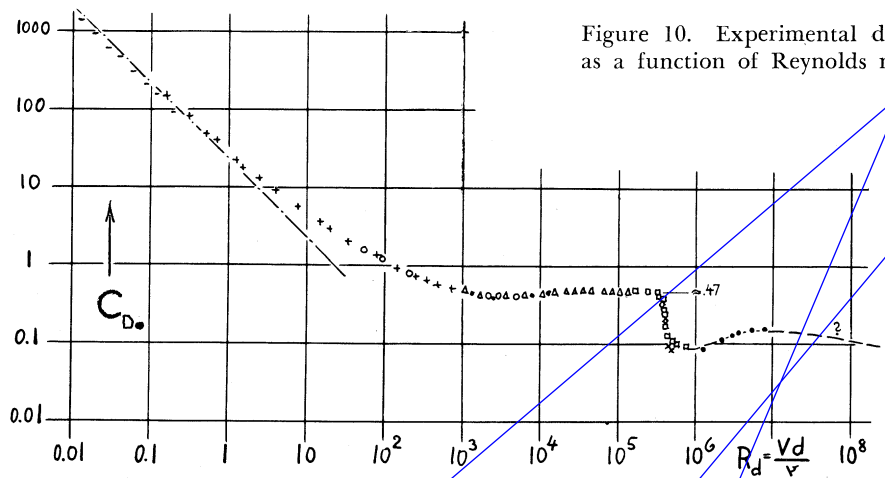

Figure 10. Experimental d
as a function of Reynolds r

The drag coefficient is roughly constant in the Reynolds-number range 10^3 to (almost) 10^6. But why does the drag coefficient drop sharply around $\mathrm{Re} \sim 10^6$? The boundary-layer picture can also help us understand this behavior. To do so, first compute Re_δ, the Reynolds number of the flow in the boundary layer. The Reynolds number is defined as

$$\mathrm{Re} = \frac{\text{typical flow speed} \times \text{distance over which the flow speed varies}}{\text{kinematic viscosity}}. \tag{7.13}$$

In the boundary layer, the flow speed varies from 0 to v, so it is comparable to v. The speed changes over the boundary-layer thickness δ. So

$$\mathrm{Re}_\delta \sim \frac{v\delta}{\nu}. \tag{7.14}$$

Because $\delta \sim r \times \mathrm{Re}^{-1/2}$,

What does this "distance over which the flow speed varies" mean?

In other words, it's the characteristic length of the object. it would be the diameter of a pipe or golf ball.

Wait, I thought it was the thickness of the boundary layer, which may not be similar to the characteristic length of the object.

In this, are we assuming typical to be the average of the flow speed? would we used the approximated average value of the flow speed?

I like that this is redefined here, it makes for easy reading!

Same sometimes i get what variables mean mixed up because so many have been introduced.

This reintroduction is very helpful and explains the graph above.

If we were being more exact, would we just take the average (ie viscosity/2)? This would still be order of viscosity though.

what do you mean by comparable?

Same order of magnitude, I think.

$$Re_\delta \sim \frac{vr}{\nu} \times Re^{-1/2}. \tag{7.15}$$

The first fraction is just the regular Reynolds number Re, so

$$Re_\delta \sim Re \times Re^{-1/2} = Re^{1/2}. \tag{7.16}$$

Flows become turbulent when $Re \sim 10^3$, so the boundary layer becomes turbulent when $Re_\delta \sim 10^3$ or $Re \sim 10^6$. Hmm! Somehow, the boundary layer's becoming turbulent reduces the drag coefficient. Now recall the interpretation of the drag coefficient as the fraction of the cross-sectional area covered by the turbulent wake. When the boundary layer becomes turbulent, it sticks to the object much better, and detaches only near the back of the object. The result is less drag!

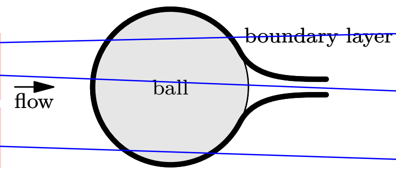

So, to get a low drag coefficient, make the object move fast enough that the Reynolds number is around 10^6. That high a Reynolds number is, however, difficult to achieve with a golf ball. That difficulty is the reason for the dimples on a golf ball. They trip the boundary layer into turbulence at a lower Reynolds number. The golf ball then travels with the benefit of this lower drag coefficient without needing to be hit at an unrealistically high speed.

Is this a general fact or did we derive it earlier in a reading? How should I know this?

Is this from the graph, since 10ˆ3 is where the curve becomes constant?

well, this is a general fact you learn in thermo, but it can also be seen in the graph. most of the flows you'll ever see are turbulent.

I feel like I've missed something here...where did this come from?

If the boundary layer becomes turbulent, is it still the boundary layer?

This seems intuitive, doesn't it?

Not to me... I'm pretty confused by it intuitively.

I think its because there is pushing and pulling from different directions, so instead of just leaving a pressure wake behind, now there is some push associated so the drag is reduced

I also don't understand this.

So are we saying that there is some force that goes against the drag force here?

This isn't intuitive to me either. And understanding this would definitely help explain why there is less drag. I understand why chain of things that happen later, but with this missing piece of understanding in the beginning makes it all kind of blurry.

$$\mathrm{Re}_\delta \sim \frac{vr}{\nu} \times \mathrm{Re}^{-1/2}. \quad (7.15)$$

The first fraction is just the regular Reynolds number Re, so

$$\mathrm{Re}_\delta \sim \mathrm{Re} \times \mathrm{Re}^{-1/2} = \mathrm{Re}^{1/2}. \quad (7.16)$$

Flows become turbulent when $\mathrm{Re} \sim 10^3$, so the boundary layer becomes turbulent when $\mathrm{Re}_\delta \sim 10^3$ or $\mathrm{Re} \sim 10^6$. Hmm! Somehow, the boundary layer's becoming turbulent reduces the drag coefficient. Now recall the interpretation of the drag coefficient as the fraction of the cross-sectional area covered by the turbulent wake. When the boundary layer becomes turbulent, it sticks to the object much better, and detaches only near the back of the object. The result is less drag!

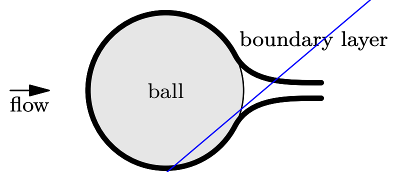

So, to get a low drag coefficient, make the object move fast enough that the Reynolds number is around 10^6. That high a Reynolds number is, however, difficult to achieve with a golf ball. That difficulty is the reason for the dimples on a golf ball. They trip the boundary layer into turbulence at a lower Reynolds number. The golf ball then travels with the benefit of this lower drag coefficient without needing to be hit at an unrealistically high speed.

How do you know it sticks better and longer?

I was wondering the same thing. Instinctively, I would think that higher turbulence would make the boundary layer more bumpy and come out the back with a larger gap, increasing the fraction of the area covered by turbulent wake. Didn't we learn in middle school/high school that turbulent flow causes more drag than laminar flow?

Is this saying that at a certain point, more turbulence is better than less turbulence?

I am also confused to why this is intuitive. A turbulent boundary layer, would seem to dettach for the object in my mind. Could we define turbulence? Maybe that is where I am going wrong.

The explanation is due to the behavior of random walks and the interpretation of viscosity as diffusion of momentum. I realize now that I should swap this and the next unit, so that we can do random walks first and then apply it to viscosity. But in lecture I'll have a go at filling in the gap.

Why does it 'stick' better?

Yeah, I'm curious about that too.

I would think that more stick would increase drag?

I'm just confused as to how turbulence would stick but reduce drag at the same time. They seem to be somewhat opposite ideas.

Think about the streamlines he drew in class. If the air doesn't stick to the object (following the edge of the object down the backside as well), a lack of air behind the object induces a decrease in pressure behind the object. This pressure different between the front and the back of the ball is what generates air drag.

The picture kind of explains it. Since the fluid sticks, the wake area is much smaller. This is because the surroundings will come around behind the ball staying close to the surface and detaches much later than it would have with less surface turbulence.

As to why it sticks with more turbulence: the more turbulence the more the surroundings slush around a given area. Picture little protrusions along a smooth surface trapping whirlpools of air.

$$\mathrm{Re}_\delta \sim \frac{\nu r}{\nu} \times \mathrm{Re}^{-1/2}. \quad (7.15)$$

The first fraction is just the regular Reynolds number Re, so

$$\mathrm{Re}_\delta \sim \mathrm{Re} \times \mathrm{Re}^{-1/2} = \mathrm{Re}^{1/2}. \quad (7.16)$$

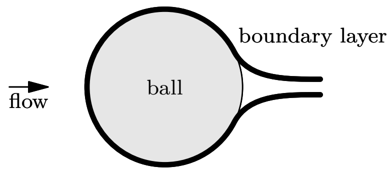

Flows become turbulent when $\mathrm{Re} \sim 10^3$, so the boundary layer becomes turbulent when $\mathrm{Re}_\delta \sim 10^3$ or $\mathrm{Re} \sim 10^6$. Hmm! Somehow, the boundary layer's becoming turbulent reduces the drag coefficient. Now recall the interpretation of the drag coefficient as the fraction of the cross-sectional area covered by the turbulent wake. When the boundary layer becomes turbulent, it sticks to the object much better, and detaches only near the back of the object. The result is less drag!

So, to get a low drag coefficient, make the object move fast enough that the Reynolds number is around 10^6. That high a Reynolds number is, however, difficult to achieve with a golf ball. That difficulty is the reason for the dimples on a golf ball. They trip the boundary layer into turbulence at a lower Reynolds number. The golf ball then travels with the benefit of this lower drag coefficient without needing to be hit at an unrealistically high speed.

Does this have anything to do with why throwing a baseball with backspin helps it carry longer in the air? the top and bottom of the ball have different turbulence due to the spin?

Backspin rotates where the flow separates and generates a vortex around the ball, just as for a wing an angle of attack generates a vortex around the wing. This vortex is the origin of lift. So the backspin generates lift (and topspin generates negative lift).

These results make sense when they are explained this way. This is a cool application.

Because of such a small r, what about other kinds of sports balls, are any others so cleverly engineered? Probably golf is the only sport where you need the ball to go so far. But like, why does a volleyball have stripes while a soccer ball (roughly the same size) has hexagonal pieces?

I think this is an interesting question! Or why a baseball has horseshoe laces or weighs 9oz...

The laces on a baseball definitely have an effect on its flight. Soccer balls and volleyball can also do strange things in the air, usually when they're hit with little spin. I think the surface differences for those may be more due to history, controllability and construction methods.

There are soccer balls that have stripes as well, I think that the design of these balls are less engineered and more cosmetic because of less technical purpose they serve. They do not need to travel as far or be as precise as a golf ball.

Agreed, I think the designs on balls like the soccer ball has more of a historical influence. The golf ball is small and generally needs to travel great distances.

Nice end summary. I am still a little bit confused about the reading, but I think reading it a couple of more times might help.

I am a little confused as well. I feel like most times when we do examples in class, we get to an answer a different way than through normal means. It seems like this was taught very similarly in 2.005.

$$\text{Re}_\delta \sim \frac{vr}{\nu} \times \text{Re}^{-1/2}. \quad (7.15)$$

The first fraction is just the regular Reynolds number Re, so

$$\text{Re}_\delta \sim \text{Re} \times \text{Re}^{-1/2} = \text{Re}^{1/2}. \quad (7.16)$$

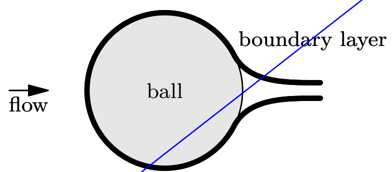

Flows become turbulent when $\text{Re} \sim 10^3$, so the boundary layer becomes turbulent when $\text{Re}_\delta \sim 10^3$ or $\text{Re} \sim 10^6$. Hmm! Somehow, the boundary layer's becoming turbulent reduces the drag coefficient. Now recall the interpretation of the drag coefficient as the fraction of the cross-sectional area covered by the turbulent wake. When the boundary layer becomes turbulent, it sticks to the object much better, and detaches only near the back of the object. The result is less drag!

So, to get a low drag coefficient, make the object move fast enough that the Reynolds number is around 10^6. That high a Reynolds number is, however, difficult to achieve with a golf ball. That difficulty is the reason for the dimples on a golf ball. They trip the boundary layer into turbulence at a lower Reynolds number. The golf ball then travels with the benefit of this lower drag coefficient without needing to be hit at an unrealistically high speed.

Do the dimples affect the ball's ability to curve? (i.e. would a non-dimpled ball slice more or less?)

The Magnus effect (the turbulent boundary layer becoming asymmetric on a spinning ball) magnifies the curve of dimpled balls.

Thanks! Any chance you'd like to clarify on why? It seems intuitively that reducing drag would reduce the tendency to curve.

This isn't technical, but I think for the same reason the ball is able to travel farther _forward_ due to the dimples, the sideways motion/ spin is also amplified by the dimples. Curve balls in baseball are also made much more effective by the laces.

a very thorough explanation. makes sense.

Cool.

Hooray! Finally this question is answered!

I wonder, how does the size of the dimples affect the Reynolds number?

I assume that the current dimple size has been thoroughly researched and studied. But would bigger dimples decrease the effectiveness? What about smaller dimples?

Maybe golf balls with dimples that are too small create a boundary layer that is not "turbulent-enough".

Quite interesting. Is this the only reason?

This is the explanation I was looking forward to the most after taking the diagnostic, and I am glad I actually understand it! :D

This was a wonderful tie in to the original problem- which we all still remember. Definitely wish there was more of a follow up on how the dimples were calculated to minimize drag coefficient

This is really cool. I remember this question from the desert island diagnostic pretest and was really confused and curious when reading the question about the golf balls, and now it is finally answered. It's really interesting to see the engineering applications in sports.

I agree and really like this example because it was hardly intuitive when i first came across it. I'm glad we revisited it in this chapter.

$$\mathrm{Re}_\delta \sim \frac{vr}{\nu} \times \mathrm{Re}^{-1/2}. \quad (7.15)$$

The first fraction is just the regular Reynolds number Re, so

$$\mathrm{Re}_\delta \sim \mathrm{Re} \times \mathrm{Re}^{-1/2} = \mathrm{Re}^{1/2}. \quad (7.16)$$

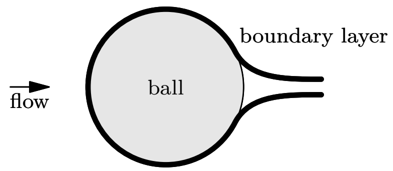

Flows become turbulent when $\mathrm{Re} \sim 10^3$, so the boundary layer becomes turbulent when $\mathrm{Re}_\delta \sim 10^3$ or $\mathrm{Re} \sim 10^6$. Hmm! Somehow, the boundary layer's becoming turbulent reduces the drag coefficient. Now recall the interpretation of the drag coefficient as the fraction of the cross-sectional area covered by the turbulent wake. When the boundary layer becomes turbulent, it sticks to the object much better, and detaches only near the back of the object. The result is less drag!

So, to get a low drag coefficient, make the object move fast enough that the Reynolds number is around 10^6. That high a Reynolds number is, however, difficult to achieve with a golf ball. That difficulty is the reason for the dimples on a golf ball. They trip the boundary layer into turbulence at a lower Reynolds number. The golf ball then travels with the benefit of this lower drag coefficient without needing to be hit at an unrealistically high speed.

I love how the questions (or variations of them) from the pre-test are answered in different sections!

How does it do this?

I think I need a bit more explanation as well. But this is really cool. I love the explanation of a nearly-everyday object!

The dimples cause the air around the golf ball to become more displaced. However when the air is interacting within a certain dimple, the Cd is low because of the high contact.

I like the explanation, but am left hanging here, so I also would like to hear a little more on this. Otherwise, good ending to the section.

In lecture I'll show pictures of the airflow with and without a turbulent boundary layer.

I wish you had talked more about this- you built it up so much in class all semester but then its all summed up in a few sentences? it's more interesting then the rest of this passage.. just an idea

Is this because it increases the distance over which the air travels

I think this is a perfect wrap up to things. I like how easily it can be placed in real world design factors.

I agree its always nice to see an actual application after the science behind it is explained.

Doesn't the turbulent flow destroy the boundary layer allowing less friction through mixing or something like that?

another diagram would be useful here (demonstrating changes in boundary layer caused by dimples?)

pretty neat.

$$\mathrm{Re}_\delta \sim \frac{v r}{\nu} \times \mathrm{Re}^{-1/2}. \tag{7.15}$$

The first fraction is just the regular Reynolds number Re, so

$$\mathrm{Re}_\delta \sim \mathrm{Re} \times \mathrm{Re}^{-1/2} = \mathrm{Re}^{1/2}. \tag{7.16}$$

Flows become turbulent when $\mathrm{Re} \sim 10^3$, so the boundary layer becomes turbulent when $\mathrm{Re}_\delta \sim 10^3$ or $\mathrm{Re} \sim 10^6$. Hmm! Somehow, the boundary layer's becoming turbulent reduces the drag coefficient. Now recall the interpretation of the drag coefficient as the fraction of the cross-sectional area covered by the turbulent wake. When the boundary layer becomes turbulent, it sticks to the object much better, and detaches only near the back of the object. The result is less drag!

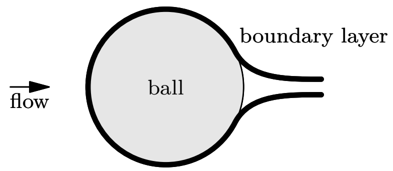

So, to get a low drag coefficient, make the object move fast enough that the Reynolds number is around 10^6. That high a Reynolds number is, however, difficult to achieve with a golf ball. That difficulty is the reason for the dimples on a golf ball. They trip the boundary layer into turbulence at a lower Reynolds number. The golf ball then travels with the benefit of this lower drag coefficient without needing to be hit at an unrealistically high speed.

Are there any other examples of things that fly farther by intentionally creating turbulence?

I think sharks can adjust the roughness of their skin in order to minimize drag – at least, that was one theory of what they were doing (and the basis of a project at Caltech to make low-drag submarines).

Is there any similar application in aviation? I know most fixed wing aircraft would stall in turbulent flow, but what about some rocket propulsion systems? Why don't rockets use this same technique?

Thats a good question. I feel like with more complex systems in aviation, this can be balanced using mechanical systems or electronics. Also the fact that most of those such systems will approach high speeds anyways, so its a little unnecessary.

this is a really nice and clear explanation of the benefits of turbulent flow and drag reduction. again the diagrams are crucial in conveying the idea.

I would just like to refer everyone to the Mythbusters episode where they took a car, covered it in clay, then carved out dimples. It had the same effect there and they showed the flow profile over it. pretty cool stuff, although they didnt go into Re #s for the episode.

I remember that episode. Didn't they not find a reasonable effect for having dimples on a car? They didn't deal with the effect of reynolds number at various speeds. For a golfball the size of a car, it should need less speed to get the same reynolds number since the size of the object increases the reynolds number already. I would think that this would make it easier for them to find an effect. I wonder if the shape of the car being different than that of a ball had an effect on their investigation.

Overall, while I understand that this section was the end of the chapter on lumping and was therefore going to be more advanced, I still found it to be an extremely technical read. It definitely took me a couple times to go over it, and I think it would be helpful for future use if it was less technical-heavy. Granted it's difficult to do that for a technical-oriented class, so I think having simpler explanations would help along the way, even if some of that information is redundant from earlier sections in the chapter.

8
Probabilistic reasoning

8.1 Is it my telephone number? 151
8.2 Why divide and conquer works 155
8.3 Characterizing distributions 167
8.4 Laplace's law of succession 168
8.5 Random walks 168

The set of tools under discussion in this part – which so far have included easy cases and lumping – all offer ways to handle complexity by discarding information. The next tool in this set is probabilistic reasoning. Instead of trying to force the world to give us exact and complete information, we accept that our information will be incomplete and use probability to quantify the incompleteness.

8.1 Is it my telephone number?

To introduce the fundamental ideas, here is an example from soon after I had moved to England sometime in the last century. I was talking to a friend on the phone, and he needed to call me back. Having just moved to the apartment, I was still unsure of my phone number (plus, British phone numbers have a strange format). I had a guess, which I thought was reasonably likely, but I was not very sure of it. To test it, I picked up my phone, dialed the candidate number, and got a busy signal.

▶ *Given this experimental evidence, how sure am I that the candidate number is my phone number? To make the question quantitative: What odds should I give?*

This question makes no sense if one uses the view of probability as long-run frequency. In that view – alas, the most common view – the probability of, say, a coin turning up heads is 1/2 because 1/2 is the limiting

do you discard information when working with easy cases?

Yes, all the information in between the easy cases! its like using boundary conditions to figure out what the middle section must be, without ever actually looking at it

Discarding? Meaning: eliminating information we have already been given to make the problem an "easy case"? I think that clears up a lot of what we have been doing lately.

Maybe you should use the phrase "lossy" since you use it everywhere else when talking about this.

I agree, it would make things more consistent.

At least one adjective in there to make sure it's clear that we're not discarding anything important!

how do you know that your information is incomplete?

This seems like an idea that we have been using all semester, even if it does fit best under "lossy" methods

I like this concept. I do feel like I am always throwing away information in this class to calculate something. It is pretty cool when the calculations end up working out.

I wouldn't say we throw out info in all the methods – we just choose numbers that are easy to deal with. In the end, we get a sense of what the answer tells us, which is much more important then a precise number you don't understand

Still, we don't have a solid basis for how close we should be to the actual result, besides (perhaps) gut instinct and some idea of how wide our estimates were along the way. This (I'm guessing) will give us a chance to quantify our wrongness.

8
Probabilistic reasoning

8.1 Is it my telephone number? 151
8.2 Why divide and conquer works 155
8.3 Characterizing distributions 167
8.4 Laplace's law of succession 168
8.5 Random walks 168

The set of tools under discussion in this part – which so far have included easy cases and lumping – all offer ways to handle complexity by discarding information. The next tool in this set is probabilistic reasoning. Instead of trying to force the world to give us exact and complete information, we accept that our information will be incomplete and use probability to quantify the incompleteness.

8.1 Is it my telephone number?

To introduce the fundamental ideas, here is an example from soon after I had moved to England sometime in the last century. I was talking to a friend on the phone, and he needed to call me back. Having just moved to the apartment, I was still unsure of my phone number (plus, British phone numbers have a strange format). I had a guess, which I thought was reasonably likely, but I was not very sure of it. To test it, I picked up my phone, dialed the candidate number, and got a busy signal.

▶ *Given this experimental evidence, how sure am I that the candidate number is my phone number? To make the question quantitative: What odds should I give?*

This question makes no sense if one uses the view of probability as long-run frequency. In that view – alas, the most common view – the probability of, say, a coin turning up heads is 1/2 because 1/2 is the limiting

Funny, I just did this not an hour or two ago to calculate my 6.055 grade so. We don't have all our memo grades for the rest of the semester, so the best estimate is an average of what we already have. Assuming we get better at writing memo's over the semester, this method would give a lower bound on our grade.

So in essence we are quantifying that "off factor" that we have been discounting through rounding up and then down throughout the semester?

I haven't read this section yet, but it would be great to finally look at these rounding errors

i disagree with the assumption that our grade will stay the same or get better. as we progress through the semester, we use up our free extensions, etc. thus, doing the same as you've always done, towards the end of the semester, will hurt you more.

Read this section for Friday's lecture (memo due Friday at 9am).

I really like this shift to am deadlines. I really have a tendency to do work later at night than the previous deadlines were set. Now its much easier to get my memos in on time

I agree. I also really liked this section title. It's quite inviting (no offense to the sections on drag...).

Assumin he has time to read before lecture I agree. If this means he doesn't or can read fewer comments, I don't mind getting my work in before 10.

Hahaha, this is a funny phrase - it deceptively makes it sound as though it were very long ago.

I like this little story

8
Probabilistic reasoning

8.1	Is it my telephone number?	151
8.2	Why divide and conquer works	155
8.3	Characterizing distributions	167
8.4	Laplace's law of succession	168
8.5	Random walks	168

The set of tools under discussion in this part – which so far have included easy cases and lumping – all offer ways to handle complexity by discarding information. The next tool in this set is probabilistic reasoning. Instead of trying to force the world to give us exact and complete information, we accept that our information will be incomplete and use probability to quantify the incompleteness.

8.1 Is it my telephone number?

To introduce the fundamental ideas, here is an example from soon after I had moved to England sometime in the last century. I was talking to a friend on the phone, and he needed to call me back. Having just moved to the apartment, I was still unsure of my phone number (plus, British phone numbers have a strange format). I had a guess, which I thought was reasonably likely, but I was not very sure of it. To test it, I picked up my phone, dialed the candidate number, and got a busy signal.

Given this experimental evidence, how sure am I that the candidate number is my phone number? To make the question quantitative: What odds should I give?

This question makes no sense if one uses the view of probability as long-run frequency. In that view – alas, the most common view – the probability of, say, a coin turning up heads is 1/2 because 1/2 is the limiting

What format could the phone numbers be in? I thought all countries had a country code (it just happens that the US has a 1) and then then after that comes an area code and the specific location digits.

Yes, but they're chunked differently. We're very well-trained to remember phone numbers as 3 digits, then 3 digits, then 4 digits. Trying to remember a different number or blocking of digits makes it much harder.

but couldn't you re-chunk it in your head? i often end up doing stuff like that.

Maybe he's referring to how it's chunked differently depending on location? http://en.wikipedia.org/wiki/Telephone_numbers_in_the_United_Kingdom

UK phone numbers, like phone numbers in many European countries, have a variable-length area code (which starts with a 0) and a variable-length number. I think it's now true that all numbers in one area have the same length, but I don't know if that was always true. So, the system makes sense but is strange to an American.

Really, they have variable length as well? That seems so...odd! Wow. They must think that our version is too simplistic when they come here.

Do you have any idea how phone numbers from instant messaging clients are generated? For instance, when I receive a text message from my friend using AIM, it's almost always 25060.

I think that it is funny that you called yourself to figure out if you remember your phone number correctly- did that really happen?

It's a true story, and is how I tested my guess. The example became Exercise 3.13 in David Mackay's textbook _Information Theory, Inference, and Learning Algorithms_ (available online too). [And It was David with whom I was talking.]

I never knew that you could call yourself. Interesting...

On cellphones...how do you think you get to your voicemail?

8
Probabilistic reasoning

8.1 Is it my telephone number? 151
8.2 Why divide and conquer works 155
8.3 Characterizing distributions 167
8.4 Laplace's law of succession 168
8.5 Random walks 168

The set of tools under discussion in this part – which so far have included easy cases and lumping – all offer ways to handle complexity by discarding information. The next tool in this set is probabilistic reasoning. Instead of trying to force the world to give us exact and complete information, we accept that our information will be incomplete and use probability to quantify the incompleteness.

8.1 Is it my telephone number?

To introduce the fundamental ideas, here is an example from soon after I had moved to England sometime in the last century. I was talking to a friend on the phone, and he needed to call me back. Having just moved to the apartment, I was still unsure of my phone number (plus, British phone numbers have a strange format). I had a guess, which I thought was reasonably likely, but I was not very sure of it. To test it, I picked up my phone, dialed the candidate number, and got a busy signal.

▶ *Given this experimental evidence, how sure am I that the candidate number is my phone number? To make the question quantitative: What odds should I give?*

This question makes no sense if one uses the view of probability as long-run frequency. In that view – alas, the most common view – the probability of, say, a coin turning up heads is 1/2 because 1/2 is the limiting

did you use a cell phone and a land line or something? How did you call with a phone that was already being used?

I was assuming you were still on the phone with your friend at the time. Maybe you should point out that you hung up with him first.

You can call your own number with a land line. Some people do it to talk to other people in different rooms in the same house. If you only have one phone connected in that house then you would probably get a busy signal but if you have multiple phones then the other phones will ring.

Didn't know you could do that.. sounds pretty lazy though.

Right, I hung up, then used the same phone. I didn't have a cell phone [and still don't have one].

Just curious- is this actually what ran through your mind after the incident?

Yes! It also helped that I had moved to England and joined the Bayesian inference group in the physics department.

Hehe cute

As a gambler, I love odds-making.

8

Probabilistic reasoning

8.1 Is it my telephone number? 151
8.2 Why divide and conquer works 155
8.3 Characterizing distributions 167
8.4 Laplace's law of succession 168
8.5 Random walks 168

The set of tools under discussion in this part – which so far have included easy cases and lumping – all offer ways to handle complexity by discarding information. The next tool in this set is probabilistic reasoning. Instead of trying to force the world to give us exact and complete information, we accept that our information will be incomplete and use probability to quantify the incompleteness.

8.1 Is it my telephone number?

To introduce the fundamental ideas, here is an example from soon after I had moved to England sometime in the last century. I was talking to a friend on the phone, and he needed to call me back. Having just moved to the apartment, I was still unsure of my phone number (plus, British phone numbers have a strange format). I had a guess, which I thought was reasonably likely, but I was not very sure of it. To test it, I picked up my phone, dialed the candidate number, and got a busy signal.

▶ *Given this experimental evidence, how sure am I that the candidate number is my phone number? To make the question quantitative: What odds should I give?*

This question makes no sense if one uses the view of probability as long-run frequency. In that view – alas, the most common view – the probability of, say, a coin turning up heads is 1/2 because 1/2 is the limiting

I'd guess about 100% chance

There's still very possible that it's some one else's busy signal though. I would guess closer to 50%.

so we're talking about the probability that you dialed a number that wasn't yours, it was busy, given that it worked (you at least got a busy signal) and you knew most of the numbers. I'd think that it was pretty unlikely that all this could happen. What if you called again in 20 mins?

I was thinking it would be 75-80% accurate. I think it's be right about 95% if you called it again in 20 mins. 99% in an hour.

I agree with your assessment on how it would go. Theoretically, would it ever be 100% sure to be your number?

Yea, like said above, I would just wait 20 mins and call again. That way if its busy again, you're almost certain that it's your number.

I trust Sanjoy's memory so I'd say 99%.

Woah this is a cool relatable example - I feel like this could very well happen to me if I moved somewhere else...

I smell bayesian probabilities coming on.

I thought this is a great introduction to the section

I agree. Not only is it something that can easily be understood, it is also something that relates to everyone.

I agree. It got me interested very quickly.

I agree too...this little story is a great way to get people engaged and interested in the new topic of probabilistic reasoning.

Are you saying you are quantifying probability?

I think it means he is using probability to quantify uncertainty.

8
Probabilistic reasoning

8.1 Is it my telephone number? 151
8.2 Why divide and conquer works 155
8.3 Characterizing distributions 167
8.4 Laplace's law of succession 168
8.5 Random walks 168

The set of tools under discussion in this part – which so far have included easy cases and lumping – all offer ways to handle complexity by discarding information. The next tool in this set is probabilistic reasoning. Instead of trying to force the world to give us exact and complete information, we accept that our information will be incomplete and use probability to quantify the incompleteness.

8.1 Is it my telephone number?

To introduce the fundamental ideas, here is an example from soon after I had moved to England sometime in the last century. I was talking to a friend on the phone, and he needed to call me back. Having just moved to the apartment, I was still unsure of my phone number (plus, British phone numbers have a strange format). I had a guess, which I thought was reasonably likely, but I was not very sure of it. To test it, I picked up my phone, dialed the candidate number, and got a busy signal.

▶ *Given this experimental evidence, how sure am I that the candidate number is my phone number? To make the question quantitative: What odds should I give?*

This question makes no sense if one uses the view of probability as long-run frequency. In that view – alas, the most common view – the probability of, say, a coin turning up heads is 1/2 because 1/2 is the limiting

In case anyone else was confused by the term "long run frequency"- it refers to the idea that as the number of trials approaches infinity, the relative frequency will converge exactly to the probability lim (x-> inf) N_x/N_t

Hmm, so is it similar to the Law of Large Numbers, then?

yes...Long-Run Frequency is another name for the Law of Large Numbers. ....an explanation of Long-Run Frequency would probably be useful in the book. (maybe in a side note?)

the relative frequency of what will converge?

Is this like considering only end-behavior?

why does it make no sense? i'm confused as to what your point is here.

You can't just look at frequencies, it wouldn't work. You need to approach it differently

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

Comments on page 2

Typo: should be evaluation.

should actually be "evaluating" i think, not "evaluation".

Right, or "evaluation of"

I feel that this sentence could be worded a little better

Yeah me too, I think it's saying that you can't extract that 1/2 from continuous flips, the 1/2 is from the probability of a single flip.

I didn't understand that's what it was trying to say until I read your comment.

Also: Rosencranz and Guildenstern are Dead.

i did a double take and now understand what you mean, but when i first read it, the word "series" triggered the idea of infinite series in my head, since i'm familiar with them for probability.

typo "evaluating"

This is interesting. I dont know much about probability, and normally when I think of probability this is what I think of...

This seems to be a very strange thing to look for since we are not given imuch information on what we want to claculate.

This reminds me of the way we have to establish that we're dealing with an "ensemble of identically-prepared systems" in thermodynamics, since we can't repeat the same process on a single system.

we could estimate the number of phone numbers in England, and the percentage of those in use at that moment

Yeah, I was going to go digit by digit (How sure are you of the first digit, how sure are you about the second, etc). Given that there are 10 choices for each digit... etc.

I feel like it would be easier just to try the same number again later, then the chance of it being busy both times if it isn't your number is way low.

I am a bit confused by what this term means can you define it

he discusses this idea in the following paragraph

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

but if you called the number over and over again, and it was always busy, wouldn't that make you more confident in your guess? or am I interpreting the frequentist interpretation incorrectly...

I really like that idea- I didn't even think of it

I agree. I would think repeating the call would be helpful in that someone may answer the number on a future call, which would prove the number was incorrect? You could even treat the person being on the line (creating the busy signal) as a fixed probability based on the average person's phone use. So I think I may have also missed something in the question.

That's a good idea. But it's not the same as repeating the same experiment many times, because with each phone call you're collecting additional evidence.

the way that I read the first full sentence here, it made it seem like frequentist interpretation would be explained in the next sentence (or paragraph), but then the question/answer that you pose after made me re-think this...I went back to the beginning of the paragraph to try and figure it out (and it took a couple of reads).

I'm still not entirely sure why it makes _no_ sense...I mean if you try enough times with reasonable timing between calls, you will reach the point where you're pretty damn sure it's the right number.

I guess what I'm saying is that i'd rather see a short sentence explaining than a question/answer that's just repeating the same thing.

couldnt you call the number multiple times? the more you call, then the problablility approaches 1 that it is your number

that's a very different situation, because you are gaining certainty in your answer (yes or no) over time. in the coin drop, your answer is a fraction, and your calculations converge to that fraction over time.

Indeed. The actual probability of the situation is either 0 or 1 - the distinction needs to be drawn between the probability that the phone is yours and how strongly you believe the number you gave is the correct number.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

I would figure it would be similar to picking a phone number at random within the same area code, potentially one in millions

This was also my first impression, but I think the fact that he wasn't completely unsure makes it difficult to quantify. For instance, I'm sure he'd know if his number was 111-1111, or something like this.

Even ignoring the simple cases for memorable phone numbers, its still a whole lot of possibilities...I feel like the work that would have to be put into eliminating obvious choices wouldn't change the number of options significantly

I agree, I think it has a lot to do with the number of potential choices presented, and how one could sort and limit them to reach a probabalistic answer.

Or is it a property of the flipper? What if the flipper was really good and could reproduce his flip of the coin with enough accurary that it became more likely for a certain side to come up. What if the flip were preformed by a robot. Then would the results of a coin toss be deterministic. this would be an interesting social experiment. Would people still consider a coin toss a random even if performed by a machine?

I'm pretty sure this is the idea behind pseudo-randomness. Most people still consider that to be "random," despite it not being by any means random, only seeming so.

The more you know about the situation, the more you know about H versus T. Friends of my parents knew Persi Diaconis well; he's a famous statistician and also a magician. I remember at dinner parties that he could make the coin come up any way he wanted. So, if he said, "I'm going to toss it in the air, it'll flip a lot, and it'll come up heads," my probability of heads would be pretty high – even though the coin was fair.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

this is going to be a tough concept to get out of my head, because that's definitely how i view coin toss probability

i don't think you should get it out of your head...it's a system that works in this situation (and works _very_ well too).

Just don't expect the method to always work.

Just think of it as: the coin is going to come up heads or it's going to come up tails. To an extent (at least in Newtonian mechanics), the outcome is predetermined – you just have no reason to suspect one outcome more than the other based on your current knowledge (for a fair coin, of course).

I am wondering why this is so. You say that the interpretation is incorrect but don't really explain why.

I'm a bit unclear if you're also saying that the placement is incorrect for the telephone example, or for the coin also. If it is also for the coin, I don't understand why either.

yeah i'm also really confused about why the placement is incorrect...I feel like a different word should be used rather than "placement" to make it less confusing.

i'm confused by this paragraph as well. i dont have a really good grasp of probability to begin with, and i dont understand the differences being outlined here.

does this mean that probability cannot be placed in the physical system itself?

I think he's saying in this type of problem you cannot.

It certainly seems right that in the phone number example the probabilities are not objective. But the text seems to endorse the stronger view that probabilities are never objective. But aren't the probabilities in quantum mechanics objective (at least, on some interpretations of quantum mechanics)? And even with the coin flip, you might think that the symmetry of the coin is an objective, physical property of the coin that can ground objective probabilities of heads and tails. Why not be a pluralist and say that some probabilities are objective and some are subjective?

y is it incorrect?

and moreover, what is then correct?

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \qquad (8.1)$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

Is this saying that the frequentist approach mutually exclusive from the Bayesian interpretation of probability? If so, if you had enough data in this specific case, wouldn't the frequentist approach apply also?

I take it Sanjoy is a Bayesian.

All of the lead-up he does to Bayes' theorem is incontrovertible - it's simple math. There's only the sense of "being a Bayesian" because the probability of E given H, for example, is rarely obvious. If you have solid rationale for the numbers you use, being a "non-Bayesian" just doesn't even make sense. Though, as with all statistics, they can be used poorly and to the benefit of whoever is paying the bill of the statistician, so you have to analyze all the assumptions being made.

Definitely!

i just wanted to say that i find this reading super interesting! :)

This is a really good layman's definition of Bayesian probability. I wish it was stated this clearly in 6.01/6.042.

typo

Might be nice here to just provide a 2X2 table to quickly compare the frequentist and Bayesian probability definitions. (2 cells would just be 'Bayesian' or 'Frequentist)

just "certainty" seems clearer than "state of certainty"

I disagree. I don't think you want to use "certainty," since that implies that you are certain about things. A "state of certainty" reflects how certain you are. It's subtly different.

OK, I'll go along with that. I'm just used to hearing certainty as "degree of certainty". "State" implies you're in or out, certain or not, which is also not what we mean. I think degree of certainty would be better here, and a comparative google search has "degree of certainty" 100x more common and more related to probability than "state of certainty".

I really like the "degree of certainty" phrase as well. It does give a feel of a measure more than "state" does.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

This reminds me of a friend of mine who likes to claim that all probability is 50%-50% – either it will happen, or it won't!

Is this directly related to the person's assumptions? therefore different assumptions result in differences in calculating assumptions

Right (if you include background knowledge as part of the assumptions)!

Could you clarify this? I would push back and say Bayesian probability is objective in the sense that you have defined formulas that you use.

But elements in Bayes' rule are subjective, like picking a prior probability.

Yes, everything hinges on the prior. In 6.041 we kind of just take the prior for granted, but in reality, all of the calculations depends on the subjectivity of the prior probability.

Does this mean that the probability of something happening can change as we gain more knowledge of it, or even that it will change?

Yes, as you know more, the probability will change. From an example I gave in an earlier comment: If you know that someone is gong to flip a coin, and that's all you know, you'll say P(heads)=1/2. But if you find out that the "someone" is a talented magician who has said, "I will flip it high in the air and it'll come up heads,", you might say P(heads)=0.9 or 0.95. (I knew someone who could do that.)

this paragraph is explained a lot better.

This sums up pretty nicely why bayesian probability is subjective. Its based on our own current beliefs.

I'd never heard this idea explained so quickly and yet make so much sense. I wish someone had said it this way earlier

It also makes the whole Monty Hall thing make more sense to me. But yeah, this is a really clear explanation.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

wait can we see an example right here before moving on to make sure we understand?

I think what he means here is that given everyone's different levels of knowledge, everyone may come up with a different probability for the same problem

How does this relate to coin flipping? Don't we all agree it is 1/2? How is that subjective?

Does this mean Bayesian inference assumes nothing probablistic can be independent?

I think the coin-flipping is used as an example because we know the outcome. Suppose another person knew something about how coins degrade that make one side significantly more probably as time approaches infinity. I feel like they would have a subjective and more accurate approximation of the probability.

This idea is still vague to me even after reading the case multiple times. Maybe its because I don't have much experience in probability.

Its sort of equivalent to the more you know the better you can assume things...like why Sanjoy can make all these really fast guesses in class that confound most of us. If you know more about the situation, assuming its more complicated than how many times a coin has been flipped, it tells you more about what you should assume the future probabilites will converge to.

I feel like the first idea is kind of in the second idea too. Collecting information is also subjective; people choose what kind of information to collect.

And there can be a bias as to what information is available to be collected. This is especially true in surveys; certain types of people are more likely to respond.

Aren't these two related? the more personal information is on an individual basis, should the range of expectation change?

yes, they are _very_ related...however they are not the _same_. 1 = it is, where 2 = it can change.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

This is a key point I think.

So if this subjective probability allows each person to create a different set of probabilities, is this due to different understandings of the world or different understanding of the evidence, or different obvs. evidence?

I believe it's due to the latter: evidence from observations.

It seems like these are the types of subconscious probabilities we run through our minds every day, interesting

I think it's a mix of all three, with the most prevalent being a different understanding of the evidence.

I really like this way of ending the paragraph. This sentence is concise and very clear, leaving the reader with the this important idea that evidence affects probability values.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

I don't really understand how this would apply to coin tossing? No matter our personal thoughts about tossing a coin, the probability of heads is always 1/2... I can see how this would apply to the phone number case because there we have to accept that the probability of that number being ours is extremely high because the other option is extremely rare... however, unless the example is something like that, I can't really grasp the application of this idea..

You bring up an interesting point.

I believe that it can be applied to the coin toss example in the following way:

Person 1 has no insight, evidence, or experience with tossing coins. He might believe that 'heads' will appear more often than tails.

Person 2 works at an arcade and spends much of his day flipping coins. In his experience, he knows that 'heads' and 'tails' appear at nearly the same frequency.

As such, "nature" has already determined the probabilities for everything in our world. It is through our experiences that we make our own interpretations that, in some easier cases, we are able to match nature's probabilities.

> "As such, "nature" has already determined the probabilities for everything in our world. It is through our experiences that we make our own interpretations that, in some easier cases, we are able to match nature's probabilities."
>
> ...well put! Thank you

The more you know about the situation, including coin tossing, the more you know about H versus T. What follows is an example repeated from an earlier comment – I wish NB had a way of connecting threads together...

Friends of my parents knew Persi Diaconis well; he's a famous statistician and also a magician. I remember at dinner parties that he could make the coin come up any way he wanted. So, if he said, "I'm going to toss it in the air, it'll flip a lot, and it'll come up heads," my probability of heads would be pretty high – even though the coin was fair.

I little table of the variables and what they mean would be helpful here.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \quad (8.1)$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

This paragraph is extremely clear and concise.. thanks!

Yeah. While I'd like to think most of us have already come across this before, it is definitely useful as a refresher.

what is this symbol? i'm not familiar with it.

It translates to "given" as explained in the rest of the sentence.

It means H given E

Yup, basically if E MUST have happened, what is the likelihood that H happened either as a result or as well?

what would you base this probability assignment off of? (I'm assuming it's a number between 0 and 1, so where on the scale would you place it?)

This is the probability we're actually interested in (the probability of our hypothesis being true given our data), so we're not going to pick a value for it. (See the last sentence of the page.)

Do you need some statement that the evidence is true?

Here, E has been observed, so it's "true" insofar as we have seen it happen.

typo Bayes'

i hate to be the grammar nazi here, but technically, it should actually be Bayes's. an s-apostrophe is used for s's that come from the plural form. an s-apostrophe-s is used for words and names that naturally end in s, such as Bayes.

Actually, there seems to be support for both ways. I've usually heard it pronounced "Bayes rule" not "Bayeses rule" (maybe the reason for the original form here?), so this consideration applies: * If the singular possessive is difficult or awkward to pronounce with an added sibilant, do not add an extra s; these exceptions are supported by The Guardian,[17] Emory University's writing center,[18] and The American Heritage Book of English Usage.[19] Such sources permit possessive singulars like these: Socrates' later suggestion; James's house, or James' house, depending on which pronunciation is intended. (caveat: from wikipedia, whose entry has "Bayes'")

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

i feel like in the phone number case it would not be as helpful to have a formula since the evidence and the initial probability are so subjective anyway

You could have a diagram of the sample space and an event to make this more clear.

what do you mean by "mental world"?

typo? hypotheses

I'm more used to seeing this probability written as the union of H and E, is the symbol that looks like an upside down u, don't know if its worth introducing though.

I'm also more used to that from 6.041, but we use the & notation in 8.044. I guess it's just a matter of convention.

yeah I think the union sign would be better since it is a more wide spread convention when talking about probability.

Conditional Probability

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

This is a great refresher for those who have seen probability before, but I wonder if this moves too quickly for someone who has never taken 6.041/18.440.

I don't think so. I felt this section was relatively self-contained. If the reader is still shaky, there are literally tons of probability books out there to which he can refer.

Is there a good illustration (marbles?) for the joint probability being equal to P(A given B) * P(B)?

I find the best way to think about it is sequentially using a tree. Also one can always go back to the coin toss example. For example, whats the probability of flipping two heads in a row. P(H on 1st toss & H on 2nd toss) = P(H on 1st) * P(H on 2nd | heads on 1st) = (1/2)*(1/2) = 1/4.

Its probably a too simple example since the probability of the 2nd head is the same no matter what was the previous toss. Its more apparent when a loaded coin is involved.

I feel like for what we are doing in this class, this concise explanation is fit... It is a nice refresher for those who have taken a course in probability and I feel like it is a clear enough summary for people who are unfamiliar with the topic. Any longer and I agree that it would take away from the point of the section.

I have never taken probability and had to re-read the paragraph but feel confident to move on.

I have never taken a probability class and I agree it is a bit much but I generally think that is ok to understand except I don't get why the two paths to H&E must be equal?

This was very difficult for me to follow. Maybe you could include a short example to illustrate this point (such as the marble example mentioned earlier).

This explanation is really helpful (especially when reflecting on old 6.041 work).

I also took 6.041 a couplte semesters ago, but it took me reading this paragraph very slowly to really understand this line of thought. I think even breaking it up into multiple paragraphs would help clarify the ideas here.

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \tag{8.1}$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

I agree it may take away from the focus of this to go on at more length about this, but I was also confused by this having never taken a course in probability...

or the quick example of diseases and positive or negative tests would provide a bit of clarity for the uniformed reader.

I have never taken a course on probability before and I admit it is a little difficult to follow.

While I think having a probability introduction or refresher for a few pages would be helpful to those students who haven't had previous exposure to probability, it would probably (no pun intended) detract from the focus in this text, and in this section in particular. I think it would be best to have an addendum or an appendix with a probability intro or refresher to those students who need it.

I do think it was a bit rushed. It's not so much that the material is too difficult or anything, but just the way that the paragraph is written seems very condensed.

It took me a while to fully understand the notation. I feel that removing P(H&E) and instead just using P(H|E) with the explanation would be better.

This notation does take a bit to understand but I think the way it's current done is important because it allows for uniformity throughout the process.

It might also help the students who have never taken 6.041 or the equivalent.

Are these two probabilities in fact always equal to each other? Because it really doesn't seem that way to me.

I feel like we could have reached this conclusion with much less explanation. I think it made it seem more complicated than it is.

I agree. P(E|H)P(H)=P(E^H)=P(H|E)P(E) and done.

For people who have seen probability in the recent past or have a very intuitive grasp of it a really quick explanation would suffice, and for those that under stand it this paragraph can easily be skipped over. I think the positives of it being here for anyone that's confused about this step outweighs the negatives of not doing so as understanding is necessary for the rest of the section.

I find it easier to see this if there is also an "=P(E,H)" here for the joint probabilit

proportion of heads in an ever-longer series of tosses. However, for evalutin the plausibility of the phone number, this interpretation – called the frequentist interpretation – makes no sense. What is the repeated experiment analogous to tossing the coin repeatedly? There is none.

The frequentist interpretation places probability in the physical system itself, as an objective property of the system. For example, the probability of 1/2 for tossing heads is seen as a property of the coin itself. That placement is incorrect and is the reason that the frequentist interpretation cannot answer the phone-number question. The sensible alternative – that probability reflects the incompleteness of our knowledge – is known as the Bayesian interpretation of probability.

The Bayesian interpretation is based on a two simple ideas. First, probabilities reflect our state of certainty about a hypothesis. Probabilities are explicitly *subjective:* Someone with a different set of knowledge will use a different set of probabilities. Second, by collecting evidence, our state of certainty changes. In other words, evidence changes our probability assignments.

In the phone-number problem, the hypothesis H is that my candidate number is correct. For this hypothesis, I have an initial or prior probability $P(H)$. After collecting the evidence E – that when I dialed this number the phone was busy – I make a new probability assignment $P(H|E)$ (the probability of the hypothesis H given the evidence E).

The recipe for using evidence to update probabilities is known as Bayes theorem. To derive it, imagine that the mental world contains only two hypothesis H and $\overline{H}$, with probabilities $P(H)$ and $P(\overline{H}) = 1 - P(H)$, and that we have collected some evidence E. Now write the joint probability $P(H\&E)$ in two different ways. $P(H\&E)$ is the probability of H being true and E occurring. That probability is, first, the product $P(H|E)P(E)$ – namely, the probability that H is true given that E occurs times the probability that E occurs. $P(H\&E)$ is, second, the product $P(E|H)P(H)$ – namely, the probability that E occurs given that H is true times the probability that H is true. These two paths to $H\&E$ must produce identical probabilities, so

$$P(E|H)P(H) = P(H|E)P(E). \quad (8.1)$$

Our goal is the updated probability of the hypothesis, namely $P(H|E)$. It is given by

Is there a proof for this?

I am always fascinated by how 6.055 incorporates so many different disciplines together.

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \quad (8.2)$$

Similarly, for the opposite hypothesis $\overline{H}$,

$$P(\overline{H}|E) = \frac{P(E|\overline{H})P(\overline{H})}{P(E)}. \quad (8.3)$$

Using the ratio $P(H|E)/P(\overline{H}|E)$, which is known as the odds, gives an even simpler formula because $P(E)$ is common to both probabilities and therefore cancels out:

$$\frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(E|H)P(H)}{P(E|\overline{H})P(\overline{H})}. \quad (8.4)$$

The right side contains the factor $P(H)/P(\overline{H})$, which is the initial odds. Using O for odds,

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})}. \quad (8.5)$$

This result is Bayes theorem (for the case of two mutually exclusive hypotheses).

In the fraction $P(E|H)/P(E|\overline{H})$, the numerator measures how well the hypothesis H explains the evidence E; the denominator measures how well the contrary hypothesis $\overline{H}$ explains the same evidence. Their ratio, known as the likelihood ratio, measures the relative value of the two hypothesis in explaining the evidence. So, Bayes theorem has the following English translation:

$$\text{updated odds} = \text{initial odds} \times \text{relative explanatory power}. \quad (8.6)$$

Let's see how this result applies to my English telephone number. Initially I was not very sure of the phone number, so $P(H)$ is perhaps 1/2 and $O(H)$ is 1. In the likelihood ratio, the numerator $P(E|H)$ is the probability of getting a busy signal given that my guess is correct (given that H is true). If my guess is correct, I'd be dialing my own phone using my phone, so I would definitely get a busy signal: $P(E|H) = 1$. The hypothesis of a correct number is a very good explanation of the data.

The trickier estimate is $P(E|\overline{H})$: the probability of getting a busy signal given that my guess is incorrect (given that $\overline{H}$ is true). If my guess is

Comments on page 3

haha ok here's what I was looking for.

I'm really glad we're getting into Bayes' theory. I had thought it would be very useful given the work we've been doing.

I like that we are doing probability, but I'm still trying to figure out how it will work in to approximations.

I completely disagree, I find myself routinely estimating probabilities in everyday life. All the little decisions we make, like whether to take the longer route with less traffic, depends on the probabilities of the situation.

While I am familiar with conditional probability, I think that a short recap on it might be helpful for other readers, particularly outside of MIT.

I disagree...I think the line "the probability that H is true given that E occurs" on the previous page is good enough

The whole paragraph before this one was a short recap

so in this problem, the opposite hypothesis is that it isn't your number correct?

I think thats the case...and then it would become the likelihood that the phone was busy given that it was NOT his phoneline

Do you think a short introduction to probability before this lecture would be helpful?

I agree. While I understand the concepts of probability presented in this section, it would be very hard to understand without some sort of background in probability.

I think that just a tiny bit more introduction to some of the terms would be good, but too much intro, and it would detract from the point, might be better as a appendix.

I think its safe to assume students in college in engineering fields will have some experience or at least an intuitive feel for basic probability, so I think the intro included is sufficient.

So maybe I'm just being too hesitant. But this still works even though P(H_bar) is = 1-P(H)?

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \tag{8.2}$$

Similarly, for the opposite hypothesis $\overline{H}$,

$$P(\overline{H}|E) = \frac{P(E|\overline{H})P(\overline{H})}{P(E)}. \tag{8.3}$$

Using the ratio $P(H|E)/P(\overline{H}|E)$, which is known as the odds, gives an even simpler formula because $P(E)$ is common to both probabilities and therefore cancels out:

$$\frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(E|H)P(H)}{P(E|\overline{H})P(\overline{H})}. \tag{8.4}$$

The right side contains the factor $P(H)/P(\overline{H})$, which is the initial odds. Using O for odds,

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})}. \tag{8.5}$$

This result is Bayes theorem (for the case of two mutually exclusive hypotheses).

In the fraction $P(E|H)/P(E|\overline{H})$, the numerator measures how well the hypothesis H explains the evidence E; the denominator measures how well the contrary hypothesis $\overline{H}$ explains the same evidence. Their ratio, known as the likelihood ratio, measures the relative value of the two hypothesis in explaining the evidence. So, Bayes theorem has the following English translation:

$$\text{updated odds} = \text{initial odds} \times \text{relative explanatory power}. \tag{8.6}$$

Let's see how this result applies to my English telephone number. Initially I was not very sure of the phone number, so $P(H)$ is perhaps 1/2 and $O(H)$ is 1. In the likelihood ratio, the numerator $P(E|H)$ is the probability of getting a busy signal given that my guess is correct (given that H is true). If my guess is correct, I'd be dialing my own phone using my phone, so I would definitely get a busy signal: $P(E|H) = 1$. The hypothesis of a correct number is a very good explanation of the data.

The trickier estimate is $P(E|\overline{H})$: the probability of getting a busy signal given that my guess is incorrect (given that $\overline{H}$ is true). If my guess is

In this reading, "odds" refers to a ratio, for example "1 to 10". However, in common vocabulary or day-to-day colloquialism, I've heard the word "odds" being used as: "The odds are 1 in a million", which would be equivalent to "1 to 999,999".

This is a slight difference that didn't get specified clearly until the last page, when you said "1 to 10". It might make a difference occasionally, but I guess usually it won't matter that much.

The day-to-day term you mentioned is a misuse of the term. When i think of odds, I think of betting, like horse racing, where they say "the odds of a horse winning is 3 to 2".

Yeah, the "odds" refers to the favorability of choosing one outcome as opposed to another (or others). Odds in favor of an event are p/(1-p).. The odds of choosing a day of the week and choosing sunday is (1/7)/ (1-1/7)= 1/6; whereas, the probability of choosing a sunday is just p= 1/7. (info from wikipedia)

Thanks, the day of the week example really helped clarify this for me.

I read most of the comments about this, but I am still confused, I would think this could use much more explanation.

I am a little worried about using this to solve problems by approximating...I'm finding it hard to keep straight in my head.

So in reading other comments on this reading so far, it seems like everyone has a different opinion of what the difference between probability and odds are. Personally, I was taught that odds is always written as a ratio of the number of times an event will happen to the number of times an event won't happen. (i.e. if the probability of an event happening is 1/3, then the odds are 1:2.) But since there seems to be so much confusion, I think a detailed explanation of the definition of odds used here would be useful.

wait maybe this won't be so foreign to me – this seems like something i learned in high school stats.

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \quad (8.2)$$

Similarly, for the opposite hypothesis $\overline{H}$,

$$P(\overline{H}|E) = \frac{P(E|\overline{H})P(\overline{H})}{P(E)}. \quad (8.3)$$

Using the ratio $P(H|E)/P(\overline{H}|E)$, which is known as the odds, gives an even simpler formula because $P(E)$ is common to both probabilities and therefore cancels out:

$$\frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(E|H)P(H)}{P(E|\overline{H})P(\overline{H})}. \quad (8.4)$$

The right side contains the factor $P(H)/P(\overline{H})$, which is the initial odds. Using O for odds,

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})}. \quad (8.5)$$

This result is Bayes theorem (for the case of two mutually exclusive hypotheses).

In the fraction $P(E|H)/P(E|\overline{H})$, the numerator measures how well the hypothesis H explains the evidence E; the denominator measures how well the contrary hypothesis $\overline{H}$ explains the same evidence. Their ratio, known as the likelihood ratio, measures the relative value of the two hypothesis in explaining the evidence. So, Bayes theorem has the following English translation:

$$\text{updated odds} = \text{initial odds} \times \text{relative explanatory power}. \quad (8.6)$$

Let's see how this result applies to my English telephone number. Initially I was not very sure of the phone number, so $P(H)$ is perhaps 1/2 and $O(H)$ is 1. In the likelihood ratio, the numerator $P(E|H)$ is the probability of getting a busy signal given that my guess is correct (given that H is true). If my guess is correct, I'd be dialing my own phone using my phone, so I would definitely get a busy signal: $P(E|H) = 1$. The hypothesis of a correct number is a very good explanation of the data.

The trickier estimate is $P(E|\overline{H})$: the probability of getting a busy signal given that my guess is incorrect (given that $\overline{H}$ is true). If my guess is

I am still a little unsure of the difference between probability and odds, as well as how we get from this equation to that.

Let p the probability that an event occurs. Then the probability that the event does not occur = 1 - p. The odds of an event occurring refers to the ratio of the probability that the event occurs divided by the probability that the event does not occur. That is, the odds = p/(1-p).

Im a little confused about just throwing in the O for odds.

I agree. To clarify why I'm confused, the odds of getting a head for a coin is 1/2. But P(H)/P(T) = 1. Maybe I have the definition wrong...

I've been using Baye's thm a lot recently- definitely have not seen it like this, or explained like this.

For some reason, I don't remember dealing with this in 6.041. Yet this form of the theorem seems very useful.

isn't it harder in this case to guess these other probabilities than just guessing the entire probability

I'm getting lost and would like to see an example soon

This is a great explanation of Bayes' theorem

I concur, I always knew it just because I knew it and never thought about the reasoning behind it like this.

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \quad (8.2)$$

Similarly, for the opposite hypothesis $\overline{H}$,

$$P(\overline{H}|E) = \frac{P(E|\overline{H})P(\overline{H})}{P(E)}. \quad (8.3)$$

Using the ratio $P(H|E)/P(\overline{H}|E)$, which is known as the odds, gives an even simpler formula because $P(E)$ is common to both probabilities and therefore cancels out:

$$\frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(E|H)P(H)}{P(E|\overline{H})P(\overline{H})}. \quad (8.4)$$

The right side contains the factor $P(H)/P(\overline{H})$, which is the initial odds. Using O for odds,

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})}. \quad (8.5)$$

This result is Bayes theorem (for the case of two mutually exclusive hypotheses).

In the fraction $P(E|H)/P(E|\overline{H})$, the numerator measures how well the hypothesis H explains the evidence E; the denominator measures how well the contrary hypothesis $\overline{H}$ explains the same evidence. Their ratio, known as the likelihood ratio, measures the relative value of the two hypothesis in explaining the evidence. So, Bayes theorem has the following English translation:

$$\text{updated odds} = \text{initial odds} \times \text{relative explanatory power}. \quad (8.6)$$

Let's see how this result applies to my English telephone number. Initially I was not very sure of the phone number, so $P(H)$ is perhaps 1/2 and $O(H)$ is 1. In the likelihood ratio, the numerator $P(E|H)$ is the probability of getting a busy signal given that my guess is correct (given that H is true). If my guess is correct, I'd be dialing my own phone using my phone, so I would definitely get a busy signal: $P(E|H) = 1$. The hypothesis of a correct number is a very good explanation of the data.

The trickier estimate is $P(E|\overline{H})$: the probability of getting a busy signal given that my guess is incorrect (given that $\overline{H}$ is true). If my guess is

This is actually quite interesting - this is different from the Bayes' theorem one normally sees: P(A|B) = P(B|A)P(A)/P(B), but they are certainly interrelated.

Actually, it's the exact same thing we see 'normally' from Bayes' theorem. If you scroll to the bottom of page 152, you'll see the equation in the form you described above. All he did here was put Bayes' theorem into a ratio of odds. They're more than interrelated, they're the same thing!

Yeah i don't think i've ever seen it in the context of odds before.

Same, I've never seen it in the context of odds, but this is a really interesting application of Bayes theorem

Agreed! I've derived it before, but never this way and never with such an obvious example

Bayes'

I've seen Bayes theorem before, but it's nice to think of it as a clear cut ratio of which hypothesis is right.

While I agree with this statement, I feel that "explains" isn't quite the proper word here. Perhaps something more like "predicts" would be better. Because fundmantally, this is the probability that E occurs given H is true, which translates to the probability that we would find evidence E if our hypothesis were true. "Explains" sounds as if our hypothesis were already true.

I agree with this. I dont really understand what "explains" would mean here. the above statement clears things up a little but its still a little fuzzy. I am not quite sure what exactly this ratio tells us in laymens' terms.

While I happen to agree with a Bayesian outlook on probabilities, I think it might be useful to mention the problem frequentists have with the Bayesian approach, namely that we need some p(H) before making any measurements, and it's hard to have a probability of a hypothesis without any evidence whatsoever.

To make it easier to follow and understand, maybe instead of separating the explanation and the example, you could merge them together.

Though, in my opinion this is one of the easier to follow and more interesting sections we have had.

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \tag{8.2}$$

Similarly, for the opposite hypothesis $\overline{H}$,

$$P(\overline{H}|E) = \frac{P(E|\overline{H})P(\overline{H})}{P(E)}. \tag{8.3}$$

Using the ratio $P(H|E)/P(\overline{H}|E)$, which is known as the odds, gives an even simpler formula because $P(E)$ is common to both probabilities and therefore cancels out:

$$\frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(E|H)P(H)}{P(E|\overline{H})P(\overline{H})}. \tag{8.4}$$

The right side contains the factor $P(H)/P(\overline{H})$, which is the initial odds. Using O for odds,

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})}. \tag{8.5}$$

This result is Bayes theorem (for the case of two mutually exclusive hypotheses).

In the fraction $P(E|H)/P(E|\overline{H})$, the numerator measures how well the hypothesis H explains the evidence E; the denominator measures how well the contrary hypothesis $\overline{H}$ explains the same evidence. Their ratio, known as the likelihood ratio, measures the relative value of the two hypothesis in explaining the evidence. So, Bayes theorem has the following English translation:

$$\text{updated odds} = \text{initial odds} \times \text{relative explanatory power}. \tag{8.6}$$

Let's see how this result applies to my English telephone number. Initially I was not very sure of the phone number, so $P(H)$ is perhaps 1/2 and $O(H)$ is 1. In the likelihood ratio, the numerator $P(E|H)$ is the probability of getting a busy signal given that my guess is correct (given that H is true). If my guess is correct, I'd be dialing my own phone using my phone, so I would definitely get a busy signal: $P(E|H) = 1$. The hypothesis of a correct number is a very good explanation of the data.

The trickier estimate is $P(E|\overline{H})$: the probability of getting a busy signal given that my guess is incorrect (given that $\overline{H}$ is true). If my guess is

i'm having trouble seeing the connection between the written words and the formulas. here i dont know what explanatory power means.

I think this is a really helpful summary–now I can go back and think through the series of formulas to make sense of everything

This is a great summary!

Wow, i never thought of conditional probability in terms of odds. This should help me remember that formula.

I agree. This word-equation is extremely helpful. And thinking about Bayes' Theorem in terms of odds is a much more intuitive way of thinking about it, at least for me.

I don't think this is as clear for people who have not taken a probability class

i really like the way this equation is set out, so that it's easy to read. thank you for the formatting (even if it is mostly LaTeX)

I agree that this way of stating the general idea of Bayes theorem is really helpful after having gone through the derivation, seeing as it gives you a way to think about it more intuitively rather than in term of the formula.

This is somewhat unclear because you have juxtaposed odds and probabilities, which on first pass, are pretty similar...

It would be nice to define odds as relative probabilities with a short equation O(H) = P(H)/P(barH)

Wouldn't O(H) be also equal to 1/2? I am confused how a probability like that can be equal to 1

Its explained later in the paragraph

I had to go back and think about this for a few seconds. Maybe spell it out even more clearly. P(H)=P(Hbar)=1/2, so 1/2 / 1/2=1 or something

I agree, it would be nice to see this more clearly listed so it could be compared more easily.

I agree with what others have been saying. It is confusing to read that O(H) is 1. What does that mean?

$$P(H|E) = \frac{P(E|H)P(H)}{P(E)}. \quad (8.2)$$

Similarly, for the opposite hypothesis $\overline{H}$,

$$P(\overline{H}|E) = \frac{P(E|\overline{H})P(\overline{H})}{P(E)}. \quad (8.3)$$

Using the ratio $P(H|E)/P(\overline{H}|E)$, which is known as the odds, gives an even simpler formula because $P(E)$ is common to both probabilities and therefore cancels out:

$$\frac{P(H|E)}{P(\overline{H}|E)} = \frac{P(E|H)P(H)}{P(E|\overline{H})P(\overline{H})}. \quad (8.4)$$

The right side contains the factor $P(H)/P(\overline{H})$, which is the initial odds. Using O for odds,

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})}. \quad (8.5)$$

This result is Bayes theorem (for the case of two mutually exclusive hypotheses).

In the fraction $P(E|H)/P(E|\overline{H})$, the numerator measures how well the hypothesis H explains the evidence E; the denominator measures how well the contrary hypothesis $\overline{H}$ explains the same evidence. Their ratio, known as the likelihood ratio, measures the relative value of the two hypothesis in explaining the evidence. So, Bayes theorem has the following English translation:

$$\text{updated odds} = \text{initial odds} \times \text{relative explanatory power}. \quad (8.6)$$

Let's see how this result applies to my English telephone number. Initially I was not very sure of the phone number, so $P(H)$ is perhaps 1/2 and $O(H)$ is 1. In the likelihood ratio, the numerator $P(E|H)$ is the probability of getting a busy signal given that my guess is correct (given that H is true). If my guess is correct, I'd be dialing my own phone using my phone, so I would definitely get a busy signal: $P(E|H) = 1$. The hypothesis of a correct number is a very good explanation of the data.

The trickier estimate is $P(E|\overline{H})$: the probability of getting a busy signal given that my guess is incorrect (given that $\overline{H}$ is true). If my guess is

Are these just guesses or is there some reasoning?

How different is the British phone system that allows you to come up with a probability of 1/2? Or did you choose 1/2 because it brought about easier numbers for estimation?

How did you come up with 1/2? This seems very arbitrary, is it just a guess? I would have thought the chance might be lower, given all of the possible ways there are to misremember something.

I've been itching for an example. Hopefully this will answer my questions

I'm a little confused on why we needed P(H)?

We needed P(H) to find O(H), which is just P(H) divided by P(H_bar)

I really like this example; I was having a hard time understanding how to implement this, but this example has given me more confidence.

I agree this example shows exactly how to approximate everyday probabilities. Without this example, I wouldn't have really known how to approach this kind of problem.

incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \tag{8.7}$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \tag{8.8}$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

Comments on page 4

Is it really a _random_ person's phone? It's just not yours. Maybe "random" should be used carefully in a section on probability.

I think it's fine to use the term "random." He's not defining any random variable here or getting more complicated, and randomness is inherent in probability. If you're really worried about it I guess you could use the term "arbitrary" in its place.

hmmm. You're probably (no pun intended) right. Since he has some idea of what the number is, the number dialed is biased to similar numbers.

P(E| H) assumes that it is not his phone. Given that it's not his phone, it is presumably a random phone (that necessarily exists, since we got a busy signal – i.e. it can't be an unused number).

I think the above argument is not that it will be some other phone, but rather that semantically, it's not random since the numbers involved probably restrict it to a particular geographic location (presumably near where he lived), etc. However, it is of little or no importance to the reading itself.

ah! I like this method, i was gonna say it's almost impossible to calculate this and I had no idea how to start it

I know, i was thinking about how many people might be on the phone at a given time, and how many possible phone numbers there are... and trying to figure out some probability that way.

This shows the simplicity in estimation that we sometimes lose track of with all of our formal methods. apply the problem to our own lives and it turns out we can come up w values that would be really hard to calculate

I'm confused about why the fraction of the time you use your phone is roughly the probability that you called someone else's busy number. Could you elaborate more on why this is a good approximation?

I think it's because for this part of the problem, he says to assume we've guessed incorrectly and dialed someone else's number. Therefore, if we assume all people use the phone for roughly the same fraction of the day,this fraction will equal the probability that any phone number we call will be busy.

incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \tag{8.7}$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \tag{8.8}$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

I've seen Bayes' theorem used in genetics problems; its really interesting to see you estimating the probability values whereas normally the majority of the problem concerns finding these values.

i wonder how cellphone usage affects this. at my house, most long conversations are held on personal cell phones, so as to not clog up the main family line.

I don't think this assumption holds across all households. I think that a more detailed way to do this calculation is to do an estimation of single households vs. family households and then break the family household down to say, an average of 4 people- 2 adults and 2 children. For the children, it might be reasonable to assume that one is old enough to use the phone, and the other isn't, but an even more detailed calculation would involve thinking about the distribution of ages of children.

This way of going about approximating the probability is very cool! It would have been near impossible to estimate the probability using amount of people, area codes, time spent on the phone etc. Just a plain 'ol people talk about 1/24 of the day makes things so simple!

This seems like an underestimate to me, if you consider businesses who use their phones constantly.

I was going to point out this very thing

This is an extremely interesting approximation that you make because you take into account some important factors that in the middle of a problem I would have never thought about (such as people being asleep at 3am and not performing the call during that time)

i'm actually very surprised you took those factors into account. since there's so many details we usually, in this class, throw away because we're only concerned with order of magnitude approximations.

Exactly what i thought of! Except I didn't realize this is considered "the odds"

Can we please stop boxing entire paragraphs? Thanks.

This was a good job of tying everything in. It makes a lot more sense after reading this.

I don't fully follow/agree with the reasoning here- I would have looked at the time of day it was rather than trying to average it over the entire day

incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \tag{8.7}$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \tag{8.8}$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

This seems awkward and too loing, but I'm not sure how to fix it.

I actually think it's helpful. I never thought to modify the guess based on when people are more likely to use their phone.

this seems like a very concise way to calculate the odds. originally I had thought about a tree diagram with number of people in a range of close area codes *land lines* usage all as the odds not in your favor. but this was well explained!

I agree! Sorry, I can't figure out how to un-question this. But it was a very clear way of figuring out the value without resorting to trees or lots of calculations!

wouldn't doing the experiment at a time like 3am make the experiment easier because not as many people are using the phone at this hour...thus raising the probability that a busy signal is yours?

I really enjoyed this section. I haven't really done anything with probabilities since 6.041 so it's nice to see this knowledge to good use.

But why didn't we take into account the # of other random people? It seems a lot more likely that you would get a busy signal not your own because there are so many other possible people you could have called.

I think this is looking at the probability of any busy signal, not just from calling your own phone.

Could you also frame this in terms of a probability? I naturally think in probability, and a 10/11 probability is much more intuitive to me than 10:1 odds.

does that give a 1/11 chance now (.09)?

yes.

No – it gives us a 10/11 probability that it is in fact his number.

The (subjective) probability that it's not his number has been reduced to 1/11.

incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \quad (8.7)$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \quad (8.8)$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

Interesting - not sure how I feel about using "odds" instead of simple probabilities like the original Bayes' rule gives, but I guess it's easier for someone who has no probability experience to wrap their head around.

I have had little probability experience and it seems that it would be much easier to go directly to the probability of calling another person's busy phone and get the same answer. But I imagine this is a framework for more complex estimates.

I've never fully understood the idea of "odds" and have always favored simple probabilities as a more intuitive way to understand the certainty.

Aren't odds and probabilities the same thing said a different way. Why do you favor one?

I agree, i think mixing the two makes it a bit more confusing for someone who is used to using one or the other (I've been taught probabilities so it seems a bit uneccessary to head this route).

To me, probability has always been much more intuitive. It very clearly tells me how likely an event is based on how many times it would happen given 100 trials. I definitely don't have this intuition for odds, but maybe it's just because I've used them less? Do other people have a good intuition for different values of odds?

This stuff is really cool. I've always wanted to learn about probability and odds but never had time in my schedule for 6.041.

If you really like this stuff then you should definitely take a probability course at MIT before you graduate.

Definitely something I want to do as well, probability has always confused me a bit as I've never had a background in it but it's pretty amazing

I really enjoyed this reading as well. I think I may end up taking a probability class because of it.

This is a different way than you expressed odds before.

incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \tag{8.7}$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \tag{8.8}$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

Can someone explain me what "10-to-1 odds" means? I usually mix this part up (and I don't trust my gut on this one).

I think 10 to 1 odds here means there are 10 times as many chances that you'd be calling your own phone (correct phone number) than calling a random phone (incorrect phone number).

thanks for explaining that, I always get confused when talking about "odds" as well!

10 to 1 means the same as 1 out of 11.

Overall a clear summary of Bayesian statistics. I've seen it once before and never used it much, and it was quite easy to follow here. I'm now kind of excited to think about day-to-day probabilities this way.

It also explains why Bayesian statistics can be so controversial - because there is (as with much of estimation) some guesswork.

did you have time to go through all this analysis as you were trying to tell your friend your phone number? just wondering

"became" or "becomes"?

or "has become"?

Good catch.

so 10:1 is "very" plausible? I wonder what the cutoff is....its a huge increase over the initial guess but 10:1 doesn't seem surefire to me!

This reading was very easy to follow and the content was very interesting.

were you right?

I felt like the initial question got convoluted in the explanation. I got a little confused on how the number being busy correlated to it actually being the right number. Maybe because I don't have much probability experience.

no, the initial question was "how certain am i that getting the busy signal means i got the right number" and the conclusion is, if i get a busy tone, i'm 10x more likely to have gotten the number right

incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \tag{8.7}$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \tag{8.8}$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

i don't really understand how your level of certainty in your guess didn't play a factor at all in our calculations.

usually it doesn't, it's just kind of an afterthought

In the end, this example was not basic enough to explain what issues I have. Although I can memorize a formula, it is difficult to master one. I wanted to see easy cases as to why P(E|H)P(H) = P(H|E)P(E).

That's a good idea to include in the reading. I'll give an example in class.

i agree. i still dont feel very confident about using this.

just curious, was the number that you dialed indeed your number?

don't leave us hanging!

Agreed. Now we're all curious.

It was right! And after that I wrote it down so I wouldn't have to guess again.

I really liked this reading. I felt like I was able to follow this reading a lot better than some of the more recent ones, particularly the ones on drag.

I too found this section much easier to follow. Granted it's an introductory section, I like that it didn't require much prior knowledge to grasp the example.

Yeah I agree. I'm excited to see how we end up using this. Hopefully I'll understand this section better since the introduction was easier for me to follow.

I agree, it definitely helps when the introduction to a new section is very clear, easy to follow and interesting. I leaves the reader eager to jump into the next section!

I agree. This section was very clearly presented and the example was simple and illustrative.

I also agree - this had a good balance of math and variables and explanation .

It's also easier to get a grasp of coin tosses, probability, and phones than drag coefficients and viscosity.

incorrect, I'd be dialing a random person's phone. What is the probability that a random phone is busy? I figure it's similar to the fraction of the day that my phone is busy. In my household, the adults use the phone maybe 1 hour per day (and the children are not yet able to use a phone). So the busy fraction is perhaps the ratio of 1 hour to 24 hours, or 1/24. But that's an underestimate. At 3am I would not do the experiment – in case I am wrong and wake someone up. Equally, I am not often on the phone at 3am. A more reasonable denominator is probably 10 or 12 hours, making the busy fraction roughly 0.1. In other words, $P(E|\overline{H}) \sim 0.1$. The hypothesis of an incorrect number is not a very good explanation for the data. The relative explanatory power, which is the likelihood ratio, is

$$\frac{P(E|H)}{P(E|\overline{H})} \sim \frac{1}{0.1} = 10. \tag{8.7}$$

Therefore, the updated odds are

$$O(H|E) = O(H) \times \frac{P(E|H)}{P(E|\overline{H})} = 10, \tag{8.8}$$

or 10-to-1 odds in favor of the number (at the start the odds were 1 to 1). The guess become very plausible!

Would we get examples for a case with more than just 2 options? How would that play out.

We'll do a famous example in lecture (Monty Hall).

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

COMMENTS ON PAGE 1

Another thought I had was... I do understand that dividing will eliminate errors.. but if you think about it... having many steps might exponentially increase the chances of error as well?

It seems like this should have been earlier - maybe right after divide and conquer was described

I agree this section would have helped reaffirm the power of divide and conquer and would have been helpful earlier in the course. That said, you'd probably have to move the probability section up with it, which may effect the flow of the course material.

Read this section for Monday (memo due Monday at 9am). It applies the probability tools to analyze one of the other tools (divide and conquer).

It doesn't seem like something that could be inaccurate – breaking up a big problem into smaller problems seems like such a classic problem-solving method, how could it be wrong?

I agree. Isn't this the basis of all problem solving techniques? Start with what you know in order to get to an unknown value?

Is there a reason why this section comes way after the introduction of divide-and-conquer

is this comparing to the confidence of just throwing a random number out there vs breaking it down using divide and conquer?

i've always been kinda puzzled by this. in divide-and-conquer reasoning, we guess more numbers. if we guess a lot of individual numbers, as is the case for divide and conquer, how does that not decrease our confidence?

I think in general, breaking the problem down into many smaller pieces allows us to reduce our overall error. For example, we may have generated a lot more error in the bandwidth over the atlantic example if we had simply guessed at the number of CDs that could fit on the 747. Instead we break the estimation down into more manageable components like volume/mass that we have more confidence in.

I tend to purposely overestimate some numbers and underestimate others so that the error ends up balancing, I think he's mentioned this too.

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\mathrm{cm}^2$ and less than $10^5\,\mathrm{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\mathrm{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\mathrm{cm}^2$ to $3000\,\mathrm{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\mathrm{cm} \times 27.94\,\mathrm{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

isn't it kinda of the idea which diversifying your assets? you decreases the variance overall. Kinda like the idea that it's unlikely that all of the assets will have negative gains and so some things will offset others

How well does confidence correlate with accuracy? Sometimes I feel confident about an answer, but it turns out to be way off, and vice versa.

Although this is true. I think confidence is huge in approximating. Talking to the gut!

I concur, especially since this whole process is really just an extremely educated guess, being condifent in your answer, that is, increasing your certainty, is very important.

If this section goes on to discuss estimation error, etc., then I've been waiting for this section for a long time. I think it needs to come much earlier, particularly since we've been entering meaningless ranges into the homework solutions for the whole term.

I agree, it definitely would have been nice if it came earlier.

wait. why are they meaningless ranges? granted, i haven't always gotten the right answer, but when i was confident, i always thought my answers made sense.

Is this almost going to turn into a when to lump and when not to lump to increase accuracy?

I had the same thought. That's obviously a very important step which we have often asked Sanjoy about. Maybe this will lead us to a more quantitative version of following our gut.

I had no idea that telegraphic could mean concise.

I don't think it necessarily means concise but instead visual. This method gives us a way to break down the problem into more small problems that we can better understand.

That may be what the reading is about, but telegraphic does in fact mean concise or terse, and not visual. At least according to Merriam-Webster.

Here, I think it just means to say "in other words" or the "short answer"

I agree, the word seems very strange here

I was thinking this in the first paragraph, that it must be a factor of how many branches we have, but that connecting the branches makes up for many estimations! Good to know I am picking up some of the intuition.

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

So the key idea is making sure that the small division calculation information is reliable. This is interesting because sometimes in divide and conquer we use quantities that are constants but have not really been proved.

That's what I was thinking. The use of "unproven" constants provides us some understanding towards a problem, even if we have little understanding of the overall concepts of the problem.

this is what i was talking about in my last comment! i guess i'd like to see more rigorous proof that the increased certainty in each of those pieces of info does indeed outweigh.

There might not even be a penalty, if we're lucky, since our estimates might balance themselves out.

True, although there are still uncertainties in all of our estimations. I'm hoping we can begin to quantify these errors that we've been minimizing through clever rounding

And the errors stack on top of each other, generally, especially if things get squared, etc.

Hold up, there could be two interpretations of the word "Penalty". He could either be saying that there's a small hassle to pay for having to go through the work of figuring out more quantities in order to be SURE about how we arrive at a number.

Or, he could be saying that there is a penalty in accuracy from introducing more quantities. Although, this second interpretation is what I thought when i read the sentence initially, I now believe the first interpretation to make more sense in this context.

I also read both interpretations when I started, but I think he means the second one since what we were talking about was accuracy

I agree. I think the important idea here is that although we may incur a small penalty in terms of the overall accuracy of our results, we gain in terms of understanding with more certainty the exact amount of inaccuracy we're dealing with. Therefore, we end up knowing how far off we are.

While this is true....isn't there a limit to the extent one could divide a problem into? From stats, we know that the variance in a term increases with the amount of terms we continue to add. Wouldn't the same apply if we divide a problem into many sub divisions?

lol

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

Is your goal to expose people to different types of vocabulary also? If not, I think dropping scarcely used words actually might distract the reader from your actual message. If you want people to increase their vocabulary then this is a fine way to do it.

Yep, and it works out even better when fellow readers define the terms right here!

I agree that it's a good idea to drop scarcely used words...unless you are going to include a dictionary of scarcely used words, most people wont take the time to look up what you mean & just gloss over it w/out knowing what you mean.

?

Maybe he meant to say 'alone'?

In lecture, you talked about how you hoped that there were 2 things that we were supposed to get out of the class: tools for estimating and a better understanding of the world. After all these examples about physics and drag, figuring out the area of a sheet of paper that isn't even really used in this country seems a lot less exciting/relevant.

I'm all for the easy examples! It's a nice break from the more complicated ones about drag.

Agreed! Easy examples are really helpful because we already know about what the example should be so we have something accurate to compare our estimation with.

The physics examples demonstrate how to apply the techniques we've been learning to harder problems like we'd encounter in the real world (scientifically), whereas the easier examples demonstrate how the techniques work and build up our confidence in what we're learning.

I like seeing a broad range of examples from all areas of life because I think it makes it slightly easier to apply approximations to something we haven't encountered before just by opening up our minds to something new. Not everything involves physics, thankfully :P

I agree, I enjoy seeing these different examples. Honestly, all the readings on drag are getting annoying. I like seeing how everything we're learning can be applied to many areas. Also, as a course 6 person, I'm never going to apply any of the information about drag in the real world, so it's just as useful in the real world as the sheet of paper is.

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

That's about 8.3"x 11.7". Isn't this also the legal-sized paper? You could refer to it as that, and people might be more familiar with it.

Whoops, I always thought A4 was 8.5 by 11"... turns out I was wrong :P

Why not just use the standard American paper, if you are teaching it in the US a lot of people will have no experience with European paper and don't know how close these estimations are... although if you said "standard 8 1/2 x 11 sheet of paper that would defeat the purpose. Maybe you could do something besides a piece of paper.

agreed. I have no idea what this size of paper would look like

Maybe because the European paper sizes are all related? A4 is half the size of A3 which is half the size of A2, etc. It is also the international standard of paper, so using 8.5x11 paper is like using imperial units when dealing with science that's done in SI. Also, the ISO standard papers have the nice property such that the ratios of side lengths are 1:sqrt(2).

That's a good explanation.

Maybe because he's more familiar with European standards? I do agree that, since this is written for an American audience, the readings should be more targeted.

I like familiar things as examples, new for problem sets.

I think you should consider a diagram giving the size of A4 paper in comparison to paper that your reader is likely to be familiar with. I know I just went to go look it up.

The problem with using US paper is that I know the size exactly (because it's called "8.5 by 11"), and probably everyone else does too, so it's not good for illustrating how to combine uncertainties.

Also, since we know that 8.5x11 is the standard here, there isn't much reason to estimate. You have the dimensions.

I know others have commented on this, but I don't know what A4 is....even w/out saying it's actual area, but just comparing it to the size of an 8.5 x 11 sheet would be helpful.

I realize now you do this in a few paragraphs, maybe just a tad earlier?

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

Why? If you can estimate (or know off hand) the dimensions it's simple...

yeah i agree, you can just use your hand to get a rough estimate

i think what we can do is a. easily guesstimate the side lengths and b. easily multiply two numbers. but how many of us are actually getting an idea about the *area*?

i think what we can do is a. easily guesstimate the side lengths and b. easily multiply two numbers. but how many of us are actually getting an idea about the *area*?

I agree... the following sentences prove the difficulty in guessing the area without multiplying dimensions through. It is much harder to guess the number of square cm's than to break it into lengths and widths and multiply....the beauty of divide and conquer.

Haha, this exaggeration helps drive home the point. You could also forgo the exponent and just type out the zero's for added emphasis.

i dont think this is a very useful example as there are many more sensible ways to go about it.

He's just walking us through it. I think maybe something that we could find in our daily lives (volume of an average cup) would be more useful, but he's doing the best he can given that we don't know what A4 looks like

draw

or "I drew"

Is this drawing also on estimation or something that is precisely 1 cm^2

Does it really matter? Most people are familiar with the length of 1cm, so a ballpark drawing of a square does provide some insight on the A4 estimation.

This doesnt seem the most practical way to start

that would be cumbersome tho

that would be cumbersome tho

I agree. This seems like a very inefficient way of going about it.

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than 1 cm^2 and less than 10^5 cm^2. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly 1 cm^2 and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer 300 cm^2 to 3000 cm^2 as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are 21.59 cm × 27.94 cm. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

This seems pretty arbitrary

This sentence really doesn't tell you anything more than what you've already said. redundancy?

But it does tell you more. We have effectively reduced our range of 1-10000cm^2 to 300-3000cm^2, which is substantially smaller (and more useful) than our first range.

I agree–this isn't exactly redundant. we just narrowed down our range so that now it spans 1 magnitude of 10 instead of 4

I think they're talking about the redundancy of explaining that a "few" = 3. I like it though, 300 to 3000 gives me a better picture in my head for some reason.

I like how we can narrow down our range so quickly through a quick visualization

Yeah, this is good and intuitive.

yeah this is a really cool method that I wouldn't have thought of!

even with just the above information. i begin thinking about a way to divide and conquer

I don't think this is the best word here.

Maybe accurate or...calculated? Useful?

I think it's alright. I've definitely heard British people use it this way before.

again. we come across the question of what the target audience for this book is :)

I often don't know any more which usages are American and which are British. That's maybe what happens when you mix two similar but not identical languages. "Two countries divided by a common language."

This sounds awkward, like I should be dividing the area into two parts, not breaking down the components of area.

I find this sentence confusing word-wise.

Yeah I think this is actually a grammatical mistake. There's definitely a better way to say this.

I read it a few times to try to get whats going on here.

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

This is pretty simple.

This is a good point.. I think when asked to estimate the area of something I would go directly to this approach and not even consider the approach discussed in the previous paragraph.

This sentence sounds awkward.

Yea, what?

I'm starting to see how this ties to probability. Sounds a lot like priors.

This definitely makes sense because if I look at a sheet of paper, it is much harder to guess the area of the paper than to just guess its length and/or width.

Is there a general rule for this? Is it cultural or in general will areas be easier to estimate than volumes?

I feel this depends on the situation. Volume deals with an extra dimension, so our accuracy will most likely drop with another estimation of a length. But say you examine a cylinder of water. Depending on size, you could compare the cylinder to 2L bottle of soda instead of trying to find the surface area.

Is it more experience? Or is eyeballing length just more intuitive anyways since area is a factor of lengths?

probably eyeballing lengths is more intuitive. if I think of area, I sort of imagine lengths and multiply them together to imagine the area; I can't just picture some arbitrary sq ft in my head.

I think the explanation and reference here to comfortable measurements is beneficial to understanding how your measurements were made.

I agree, we used this length measurement all semester to understand cross sectional area when looking at drag.

thats funny to think about- but makes so much sense

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\mathrm{cm}^2$ and less than $10^5\,\mathrm{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\mathrm{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\mathrm{cm}^2$ to $3000\,\mathrm{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\mathrm{cm} \times 27.94\,\mathrm{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

Hmm..that is a really great way to prove this point!

I agree, it really highlights our intuition and experience.

yeah I agree, really good way to make the reader realize that they know more about lengths than areas.

I think this is not actually a great example, because in this case a person's height is a much more useful value than his cross-sectional area.

I think a better example (although perhaps way to close to the actual problem) would be the area of a computer screen. I know my screen has a 15" diagonal, but no immediate guess about its area (which is arguably the important quantity for content viewing).

Perhaps even better would be: what's the area of a computer desk, or a dining room table, or something similar. In these cases area is (usually) an important factor, but we (or I, at least) have a way better feel for the length scale.

Wouldn't it be funny if we actually measured people in cross sectional area. Growth charts would look a little different lol

agreed, area is not a familiar quantity

I like this - it is a good simple way to break down why divide and conquer seems easier.

what point are you trying to make here? i'm confused about where this is going.

He's trying to explain why divide and conquer makes the problem easier. It's easier to estimate height and width than area because we know those numbers better.

He's justifying why its easier to estimate a length as opposed to an area.

How we think (to find area using length and height rather than just guessing area) is more intuitive, so this is how we should go about estimating an unknown when the quantity we are looking for is area

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

Very true, although not something I think about too much. I like the strategy of using estimations we are accustomed to.

I think though, this strategy is hardly worth mentioning, since it is instinctive already for us to resort to familiar estimations than unfamiliar ones. Stating the obvious?

Obvious to some yes, but unlikely obvious to everyone.

It also makes sense to include it in our current discussion as it is a good example.

Yeah, I don't think it's the strategy he's emphasizing here–we learned divide and conquer in our first unit. here he's talking about why divide and conquer produces more accurate results than straight guessing–because we change the problem to use values we are more familiar with

Yeah I think it's really interesting to think about this. It wasn't so obvious to me before, but now it makes complete sense.

I agree it's instinctive, but that most likely means we thought about why we did it less.

I think this statement is interesting, but I'd like to see something to back it up. (I feel like it's intuitive, but it would still be nice to see just by how much the difference is, or how many people, etc)

I think it would be important to note the relative accuracy of knowing each of these quantities. Saying something along the lines of 'we are far more accurate estimating the length and the width than the area, granted the higher variance in our answer for estimating two quantities instead of one.'

Doesn't this fact kind of nullify the process of approximating? It's kinda like looking up an answer to a similar problem....

Yeah I saw this and I was like...huh? Isn't that the answer to our question?

No, this isn't saying that A4 IS letter patter, it's saying that it's the equivalent standard for them. They are different sizes.

Exactly - its in the same class of paper and therefore a good place to start.

I don't really understand the point of this paragraph. It gives too many facts that seem somewhat irrelevant to our point.

8.2 Why divide and conquer works

How does divide-and-conquer reasoning (Chapter 1) produce such accurate estimates? Alas, this problem is hard to analyze directly because we do not know the accuracy in advance. But we can analyze a related problem: how divide-and-conquer reasoning increases our confidence in an estimate or, more precisely, decreases our uncertainty.

The telegraphic answer is that it works by subdividing a quantity about which we know little into several quantities about which we know more. Even if we need many subdivisions before we reach reliable information, the increased certainty outweighs the small penalty for combining many quantities. To explain that telegraphic answer, let's analyze in slow motion a short estimation problem using divide-and-conquer done.

8.2.1 Area of a sheet of paper

The slow-motion problem is to estimate area of a sheet of A4 paper (A4 is the standard sheet size in Europe). On first thought, even looking at a sheet I have little idea about its area. On second thought, I know something. For example, the area is more than $1\,\text{cm}^2$ and less than $10^5\,\text{cm}^2$. That wide range makes it hard to be wrong but is also too wide to be useful. To narrow the range, I'll drew a small square with an area of roughly $1\,\text{cm}^2$ and guess how many squares fit on the sheet: probably at least a few hundred and at most a few thousand. Turning 'few' into 3, I offer $300\,\text{cm}^2$ to $3000\,\text{cm}^2$ as a plausible range for the area.

Now let's use divide-and-conquer and get a more studied range. Subdivide the area into the width and height; about two quantities my knowledge is more precise than it is about area itself. The extra precision has a general reason and a reason specific to this problem. The general reason is that we have more experience with lengths than areas: Which is the more familiar quantity, your height or your cross-sectional area? Therefore, our length estimates are usually more accurate than our area estimates.

The reason specific to this problem is that A4 paper is the European equivalent of standard American paper. American paper is known to computers and laser printers as 'letter' paper and known commonly in the United States as 'eight-and-a-half by eleven' (inches!). In metric units, those dimensions are $21.59\,\text{cm} \times 27.94\,\text{cm}$. If A4 paper were identical to letter paper, I could now compute its exact area. However, A4 paper is, I

my computer used to have a chinese windows os and run chinese word, and i basically had these numbers memorized so i could get my pages to print correctly. just a fun fact! :)

This sentence reads awkwardly

remember from living in England, slightly thinner and longer than letter paper. I forget the exact differences between the dimensions of A4 and letter paper, hence the remaining uncertainty: I'll guess that the width lies in the range 19...21 cm and the length lies in the range 28...32 cm.

The next problem is to combine the plausible ranges for the height and width into the plausible range for the area. A first guess, because the area is the product of the width and height, is to multiply the endpoints of the width and height ranges:

$$A_{\text{min}} = 19\,\text{cm} \times 28\,\text{cm} = 532\,\text{cm}^2;$$
$$A_{\text{max}} = 21\,\text{cm} \times 32\,\text{cm} = 672\,\text{cm}^2.$$

This method turns out to overestimate the range – a mistake that I correct later – but even the too-large range spans only a factor of 1.26 whereas the starting range of 300...3000 cm^2 spans a factor of 10. Divide and conquer has significantly narrowed the range by replacing quantities about which we have little knowledge, such as the area, with quantities about which we have more knowledge.

The second bonus, which I now quantify correctly, is that subdividing into many quantities carries only a small penalty, smaller than suggested by naively multiplying endpoints. The naive method overestimates the range because it assumes the worst. To see how, imagine an extreme case: estimating a quantity that is the product of ten factors, each that you know to within a factor of 2 (in other words, each plausible range is a factor of 4). Is your plausible range for the final quantity a factor of $4^{10} \approx 10^6$?! That conclusion is terribly pessimistic. A more likely result is that a few of the ten estimates will be too large and a few too small, allow several errors to cancel.

To quantify and fix this pessimism, I will explain plausible ranges using probabilities. Probabilities are the tool for this purpose. As discussed in Section 8.1, probabilities reflect incomplete knowledge; they are *not* frequencies in a random experiment (Jaynes's *Probability Theory: The Logic of Science* [11] is a excellent, book-length discussion and application of this fundamental point).

To make a probabilistic description, start with the proposition or hypothesis

$$H \equiv \text{The area of A4 lies in the range } 300\ldots3000\,\text{cm}^2.$$

For someone who has never seen or heard of A4 paper before, though, I would have no idea how to go about estimating its size.

Agreed. I think this description of the paper would be more useful when A4 paper is first introduced. While I still don't know what it looks like, I now have an idea of it, since its described as being slightly thinner and longer.

Couldn't we just say these probably cancel out?

It's _narrower_, not thinner, right?

Haha yeah, I started thinking that he was referring to thickness too and wondered why that was pertinent to anything

is it like legal paper?

yeah when I first arrived in the US, I didn't realize that the paper was different but then I went home and brought my notebooks with me I noticed that they were really different!

seems like these two would cancel out...so you could just compute the area of letter paper and it would probably be the same as A4 area.

I used to live where A4 papers are used. It was weird when I moved here.

i really like what you're trying to do in this section, but it reads oddly to me because it sounds like sharing your own experiences. it's all based on what you happened to remember from living abroad (longer and thinner) and what you don't remember (the exact dimensions).

While we certainly don't share these exact experiences, I think it's helpful to see how we can apply our experience and intuition, even if it's not specific, to the problem at hand.

i like this use of ranges, I wish we had learned this earlier. I never used the range entry boxes in the psets

Also, you multiply the uncertainty of length and width, to get an even bigger uncertainty in the area.

I think that the purpose of the range is to account for uncertainty. Instead of choosing 20 and 30 he chose numbers over a range to allow for the chance that he's off by a couple of centimeters.

remember from living in England, slightly thinner and longer than letter paper. I forget the exact differences between the dimensions of A4 and letter paper, hence the remaining uncertainty: I'll guess that the width lies in the range 19 . . . 21 cm and the length lies in the range 28 . . . 32 cm.

The next problem is to combine the plausible ranges for the height and width into the plausible range for the area. A first guess, because the area is the product of the width and height, is to multiply the endpoints of the width and height ranges:

$$A_{\text{min}} = 19\,\text{cm} \times 28\,\text{cm} = 532\,\text{cm}^2;$$
$$A_{\text{max}} = 21\,\text{cm} \times 32\,\text{cm} = 672\,\text{cm}^2.$$

This method turns out to overestimate the range – a mistake that I correct later – but even the too-large range spans only a factor of 1.26 whereas the starting range of $300 \ldots 3000\,\text{cm}^2$ spans a factor of 10. Divide and conquer has significantly narrowed the range by replacing quantities about which we have little knowledge, such as the area, with quantities about which we have more knowledge.

The second bonus, which I now quantify correctly, is that subdividing into many quantities carries only a small penalty, smaller than suggested by naively multiplying endpoints. The naive method overestimates the range because it assumes the worst. To see how, imagine an extreme case: estimating a quantity that is the product of ten factors, each that you know to within a factor of 2 (in other words, each plausible range is a factor of 4). Is your plausible range for the final quantity a factor of $4^{10} \approx 10^6$?! That conclusion is terribly pessimistic. A more likely result is that a few of the ten estimates will be too large and a few too small, allow several errors to cancel.

To quantify and fix this pessimism, I will explain plausible ranges using probabilities. Probabilities are the tool for this purpose. As discussed in Section 8.1, probabilities reflect incomplete knowledge; they are *not* frequencies in a random experiment (Jaynes's *Probability Theory: The Logic of Science* [11] is a excellent, book-length discussion and application of this fundamental point).

To make a probabilistic description, start with the proposition or hypothesis

$$H \equiv \text{The area of A4 lies in the range } 300 \ldots 3000\,\text{cm}^2.$$

are you saying both Amin and Amax overestimate the range or just Amax

I think it means overestimate the range in the full answer 532cm^2 - 672cm^2. Each individual answer is an underestimate or overestimate respectively, and together they overestimate the range.

That's right – the range is too wide on both ends. In other words, you know more than you think you do based on that wide range.

Why would it overestimate the range?

Because we multiplied the smallest width and height to get the bottom range and the largest width and height to get the upper.

Yup. Realistically its not going to be super-skinny and super-tall, it would either be super-skinny medium tall, or super-tally medium skinny.

Nice way of explaining the reason for overestimation... thanks!

It makes sense that dividing will decrease chances of error.

I think this is a powerful statement with everyday problems. Estimations require order of magnitude guesses but usually our intuition is much better

that really is an amazing improvement

was this really divide and conquer? in both cases you 'divided' the problem into something you knew.

i thought that was the idea of divide and conquer.

That's the definition of divide and conquer though, isn't it? You divide something into something you know how to deal with.

Is narrowing the range so that it does not include the actual answer really beneficial though?

Yes we have knowledge about the hieght and length, but I had no knowledge about the relationship between A4 and 8X11

remember from living in England, slightly thinner and longer than letter paper. I forget the exact differences between the dimensions of A4 and letter paper, hence the remaining uncertainty: I'll guess that the width lies in the range 19...21 cm and the length lies in the range 28...32 cm.

The next problem is to combine the plausible ranges for the height and width into the plausible range for the area. A first guess, because the area is the product of the width and height, is to multiply the endpoints of the width and height ranges:

$$A_{\text{min}} = 19\,\text{cm} \times 28\,\text{cm} = 532\,\text{cm}^2;$$
$$A_{\text{max}} = 21\,\text{cm} \times 32\,\text{cm} = 672\,\text{cm}^2.$$

This method turns out to overestimate the range – a mistake that I correct later – but even the too-large range spans only a factor of 1.26 whereas the starting range of 300...3000 cm^2 spans a factor of 10. Divide and conquer has significantly narrowed the range by replacing quantities about which we have little knowledge, such as the area, with quantities about which we have more knowledge.

The second bonus, which I now quantify correctly, is that subdividing into many quantities carries only a small penalty, smaller than suggested by naively multiplying endpoints. The naive method overestimates the range because it assumes the worst. To see how, imagine an extreme case: estimating a quantity that is the product of ten factors, each that you know to within a factor of 2 (in other words, each plausible range is a factor of 4). Is your plausible range for the final quantity a factor of $4^{10} \approx 10^6$?! That conclusion is terribly pessimistic. A more likely result is that a few of the ten estimates will be too large and a few too small, allow several errors to cancel.

To quantify and fix this pessimism, I will explain plausible ranges using probabilities. Probabilities are the tool for this purpose. As discussed in Section 8.1, probabilities reflect incomplete knowledge; they are *not* frequencies in a random experiment (Jaynes's *Probability Theory: The Logic of Science* [11] is a excellent, book-length discussion and application of this fundamental point).

To make a probabilistic description, start with the proposition or hypothesis

$$H \equiv \text{The area of A4 lies in the range } 300\ldots3000\,\text{cm}^2.$$

So divide and conquer is really divide until you know a quantity then combine!

the combining is the conquering part

This is the evidence of our probability unit, isn't it? This is slowly making a lot of sense in terms of the new unit we're starting.

If each of the smaller sections contains smaller penalties, wouldn't they combine to result in a similar level of penalty?

I like this example...after rereading...it helps quantify why more estimates should reduce the certainty.

Typo: allowing.

remember from living in England, slightly thinner and longer than letter paper. I forget the exact differences between the dimensions of A4 and letter paper, hence the remaining uncertainty: I'll guess that the width lies in the range 19 . . . 21 cm and the length lies in the range 28 . . . 32 cm.

The next problem is to combine the plausible ranges for the height and width into the plausible range for the area. A first guess, because the area is the product of the width and height, is to multiply the endpoints of the width and height ranges:

$$A_{\text{min}} = 19\,\text{cm} \times 28\,\text{cm} = 532\,\text{cm}^2;$$
$$A_{\text{max}} = 21\,\text{cm} \times 32\,\text{cm} = 672\,\text{cm}^2.$$

This method turns out to overestimate the range – a mistake that I correct later – but even the too-large range spans only a factor of 1.26 whereas the starting range of 300 . . . 3000 cm^2 spans a factor of 10. Divide and conquer has significantly narrowed the range by replacing quantities about which we have little knowledge, such as the area, with quantities about which we have more knowledge.

The second bonus, which I now quantify correctly, is that subdividing into many quantities carries only a small penalty, smaller than suggested by naively multiplying endpoints. The naive method overestimates the range because it assumes the worst. To see how, imagine an extreme case: estimating a quantity that is the product of ten factors, each that you know to within a factor of 2 (in other words, each plausible range is a factor of 4). Is your plausible range for the final quantity a factor of $4^{10} \approx 10^6$?! That conclusion is terribly pessimistic. A more likely result is that a few of the ten estimates will be too large and a few too small, allow several errors to cancel.

To quantify and fix this pessimism, I will explain plausible ranges using probabilities. Probabilities are the tool for this purpose. As discussed in Section 8.1, probabilities reflect incomplete knowledge; they are *not* frequencies in a random experiment (Jaynes's *Probability Theory: The Logic of Science* [11] is a excellent, book-length discussion and application of this fundamental point).

To make a probabilistic description, start with the proposition or hypothesis

$$H \equiv \text{The area of A4 lies in the range } 300 \ldots 3000\,\text{cm}^2.$$

Is this something we have been relying on all semester? I really like this analysis

I remember using this for the very first pset...for example with the imported oil in the US, we knew that if we over estimated cars, we might want to fudge another factor. Though this reading brings up the great point that when you have more than a few factors that are estimated, their is no need to consciously fudge! It is more likely that it will just work out!

Yeah I feel like if you go with what you think is right everything tends to work out in the end :-). Of course, if your errors keep stacking, that's not a great place to be...

If you knew you overestimated cars then why wouldn't you lower your estimate instead of fudging another factor?

I completely agree with 5:07. This allowing uncertainties to cancel only works if the errors are, in fact, unknown. If we know that an error is off by a factor of 2 in a particular direction, we can't expect the other errors to cancel it out. (This is vaguely reminiscent of martingales, I feel, since we can approximate each estimation as having an expected error factor of 1.)

The variance increases, no doubt. But this assumption relies heavily on probabilities and expected values.

In response to 5:07: If I remember correctly, we used a high number for the number of cars because it was a round number and our estimated number of cars wasn't round. We then fudged the next calculation to a lower round number to make up for it. It makes the calculations easier if you use round numbers so you fudge up and down and balance it so that they cancel.

Answered my earlier question, I feel like this is something that should have been introduced much earlier, it is really valuable

I was wondering when this was going to come into play!

ooh nice transition into this unit

remember from living in England, slightly thinner and longer than letter paper. I forget the exact differences between the dimensions of A4 and letter paper, hence the remaining uncertainty: I'll guess that the width lies in the range 19 . . . 21 cm and the length lies in the range 28 . . . 32 cm.

The next problem is to combine the plausible ranges for the height and width into the plausible range for the area. A first guess, because the area is the product of the width and height, is to multiply the endpoints of the width and height ranges:

$$A_{\text{min}} = 19\,\text{cm} \times 28\,\text{cm} = 532\,\text{cm}^2;$$
$$A_{\text{max}} = 21\,\text{cm} \times 32\,\text{cm} = 672\,\text{cm}^2.$$

This method turns out to overestimate the range – a mistake that I correct later – but even the too-large range spans only a factor of 1.26 whereas the starting range of 300 . . . 3000 cm^2 spans a factor of 10. Divide and conquer has significantly narrowed the range by replacing quantities about which we have little knowledge, such as the area, with quantities about which we have more knowledge.

The second bonus, which I now quantify correctly, is that subdividing into many quantities carries only a small penalty, smaller than suggested by naively multiplying endpoints. The naive method overestimates the range because it assumes the worst. To see how, imagine an extreme case: estimating a quantity that is the product of ten factors, each that you know to within a factor of 2 (in other words, each plausible range is a factor of 4). Is your plausible range for the final quantity a factor of $4^{10} \approx 10^6$?! That conclusion is terribly pessimistic. A more likely result is that a few of the ten estimates will be too large and a few too small, allow several errors to cancel.

To quantify and fix this pessimism, I will explain plausible ranges using probabilities. Probabilities are the tool for this purpose. As discussed in Section 8.1, probabilities reflect incomplete knowledge; they are *not* frequencies in a random experiment (Jaynes's *Probability Theory: The Logic of Science* [11] is a excellent, book-length discussion and application of this fundamental point).

To make a probabilistic description, start with the proposition or hypothesis

$$H \equiv \text{The area of A4 lies in the range } 300 \ldots 3000\,\text{cm}^2.$$

I feel like I often make this mistake. I guess because there is a fine line between what we know and what has been done experimentally. Also, the way we're taught probability focuses on random experiment.

This is a hard concept to grasp; Thanks for the book note here; I think I need to read it, because like the above comment, I am consistently stuck in the random experiment mindset.

> Here's an example, which I'll discuss more in lecture, that shows why the random-experiment interpretation is lacking. What is the probability that the trillionth (or ten trillionth) digit of pi is a 7? There's only one pi, so there's no way to do the experiment, and it isn't random at all. But you can still give a reasonable answer for the probability.

i don't get this phrase

I think he's saying that probability doesn't just describe the frequency of some measured value in an experiment, but rather describes our state of knowledge. In other words, probabilities aren't just about the number of times a coin lands on heads if you flip it x number of times. But I agree that the use of the term frequencies here threw me off a bit at first.

yeah, I was confused by it also but your explanation really helped.

This needs more explaining

an

nice catch, i just glanced over it assuming it was already "an"

and information (or evidence)

$$I \equiv \text{What I know about the area } \textit{before} \text{ using divide and conquer.}$$

Now I want to find the conditional probability $P(H|I)$ – namely, the probability of H given my knowledge before trying divide and conquer. There is no known algorithm for computing a probability in such a complicated problem situation. How, for example, does one represent my state of knowledge? In these cases, the best we can do is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my intuitive knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300 \ldots 3000\,\text{cm}^2$. The surprise suggests that $P(H|I)$ is larger than 1/2. The mildness of the surprise suggests that $P(H|I)$ is not much larger than 1/2. I'll quantify it as $P(H|I) = 2/3$: I would give 2-to-1 odds that the true area is within the plausible range.

Furthermore, I'll use this probability or odds to define a plausible range: It is the range for which I think 2-to-1 odds is fair. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will later help give plausible ranges an intuitive interpretation (as a region on a log-normal distribution). That interpretation will then help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly 2/3 of the probability. With a further assumption of symmetry, the plausible range $300 \ldots 3000\,\text{cm}^2$ represents the following probabilities:

$$\begin{aligned} P(A < 300\,\text{cm}^2) &= 1/6; \\ P(300\,\text{cm} \le A \le 3000\,\text{cm}^2) &= 2/3; \\ P(A > 3000\,\text{cm}^2) &= 1/6. \end{aligned}$$

Here is the corresponding picture with width proportional to probability:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
A	$< 300\,\text{cm}^2$	$300 \ldots 3000\,\text{cm}^2$	$> 3000\,\text{cm}^2$

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are

Love how a simple notion can be so easily represented by variables!

Clever. I really like and understand how you have phrased this conditional probability

I like that you show how different estimating methods are somewhat related. Of course, this would explain why this is in Section 8 instead of Section 1.

I think that this section should have come before the last one. I think it is an easier introduction to probability.

I agree.

Yeah i was wondering about this too. We don't necessarily have finite states that we have probabilities about. I guess we could use continuous random variables?...

I wasn't quite sure about this, but read on and you'll see how he approximates it.

so does this really help us out in the end?

But can introspection give us an actual probability? I guess its almost what we've been doing on the problem sets with our ranges, but even then thats almost just guessing at how confident we are arbitrarily.

I think he is saying that probability is a way of explaining (in an objective and explicit manner) what we know and what we don't know. Talking to our gut (introspecting) is supposed to help us gain a better understanding of what we don't know.

This was my first worry at the start of this unit. I didn't think there could possibly be ways to quantify some of the reasoning we've been using, and thus, probability seemed like a waste of time. But now I see that we're using probability in a different and more useful way to enhance our methods.

my gut is confused and quiet! :(

and information (or evidence)

$I \equiv$ What I know about the area *before* using divide and conquer.

Now I want to find the conditional probability $P(H|I)$ – namely, the probability of H given my knowledge before trying divide and conquer. There is no known algorithm for computing a probability in such a complicated problem situation. How, for example, does one represent my state of knowledge? In these cases, the best we can do is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my intuitive knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300\ldots3000\,\text{cm}^2$. The surprise suggests that $P(H|I)$ is larger than 1/2. The mildness of the surprise suggests that $P(H|I)$ is not much larger than 1/2. I'll quantify it as $P(H|I) = 2/3$: I would give 2-to-1 odds that the true area is within the plausible range.

Furthermore, I'll use this probability or odds to define a plausible range: It is the range for which I think 2-to-1 odds is fair. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will later help give plausible ranges an intuitive interpretation (as a region on a log-normal distribution). That interpretation will then help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly 2/3 of the probability. With a further assumption of symmetry, the plausible range $300\ldots3000\,\text{cm}^2$ represents the following probabilities:

$$P(A < 300\,\text{cm}^2) = 1/6;$$
$$P(300\,\text{cm} \le A \le 3000\,\text{cm}^2) = 2/3;$$
$$P(A > 3000\,\text{cm}^2) = 1/6.$$

Here is the corresponding picture with width proportional to probability:

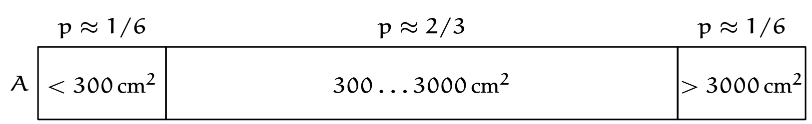

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are

I really enjoyed the class lecture about gut feelings and intuition, it would be awesome to see a chapter in this book dedicated to that idea.

I agree, I have found that having confidence in my gut and looking at things the right way is the most important part of figuring out many of these problems.

I too agree, but sometimes I wonder if I'm trusting my gut too much in some of our examples where I don't have a lot of natural intuition. Having some of these ideas as a chapter/section of the book would be helpful.

Regardless of whether it gets a whole chapter or just a mention such as this, I feel like it would be beneficial to the reader if it was mentioned earlier on, as early as you mentioned it in class.

This seems a little ironic. In this method, you are going to find the information based on what you know, but in the probabilistic description, you are going with your gut. This doesn't seem consistant.

This seems like kind of a ridiculous argument

I think it's a pretty interesting way to think about probability–how surprised would you be if your guess was wrong? as opposed to–on a scale of 0 to 1 how right do you think you are...hm interesting

for "ridiculous": Are you talking about the use of the metaphor of a guessing intestine or complaining about something more substantial? Having confidence intervals are important for gauging how to proceed, and this range is already based on some reasoning.

This is a nice paragraph. I like the idea of actually having to 'talk to the gut' in order to get an good estimate. There are times where you can just simply say "I know it's at least this" even with no basis because of the experiences I had, hence the training my 'gut' has.

Don't think of it as an argument. That is, don't think of it as a proof for the value of P(H|I). Rather, it's a measurement of my own (subjective) value of P(H|I). You can use the same instrument to measure your own P(H|I).

Sounds like a handy way to think about probability when you need to problem solve without references.

and information (or evidence)

$I \equiv$ What I know about the area *before* using divide and conquer.

Now I want to find the conditional probability $P(H|I)$ – namely, the probability of H given my knowledge before trying divide and conquer. There is no known algorithm for computing a probability in such a complicated problem situation. How, for example, does one represent my state of knowledge? In these cases, the best we can do is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my intuitive knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300 \ldots 3000\,\mathrm{cm}^2$. The surprise suggests that $P(H|I)$ is larger than 1/2. The mildness of the surprise suggests that $P(H|I)$ is not much larger than 1/2. I'll quantify it as $P(H|I) \approx 2/3$: I would give 2-to-1 odds that the true area is within the plausible range.

Furthermore, I'll use this probability or odds to define a plausible range: It is the range for which I think 2-to-1 odds is fair. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will later help give plausible ranges an intuitive interpretation (as a region on a log-normal distribution). That interpretation will then help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly 2/3 of the probability. With a further assumption of symmetry, the plausible range $300 \ldots 3000\,\mathrm{cm}^2$ represents the following probabilities:

$$P(A < 300\,\mathrm{cm}^2) = 1/6;$$
$$P(300\,\mathrm{cm} \le A \le 3000\,\mathrm{cm}^2) = 2/3;$$
$$P(A > 3000\,\mathrm{cm}^2) = 1/6.$$

Here is the corresponding picture with width proportional to probability:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
A	$< 300\,\mathrm{cm}^2$	$300 \ldots 3000\,\mathrm{cm}^2$	$> 3000\,\mathrm{cm}^2$

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are

This sounds like what I do when assigning weighted numbers for guessing on our homework sets...

That IS what we do in our homework!

I understand what you're saying here, but "talking to your gut" regarding probability can quickly lead you down a slippery slope.

I agree, in probability classes they usually tell you to not go with your gut or first instinct.

True, but you are quantifying your gut in this case, so you are going by how much you believe in your gut.

This probabilistic reasoning makes me uneasy. Many people are unlikely to intuitively understand probabilities unless they've seen them many times before. Not sure I would count on my gut in these situations.

Yeah. I also don't see how that leads to a probability of 1/2. Is this a general rule of thumb to use a half as a probability for a range that we're pretty sure about.

Yeah, I'm also feeling uneasy about using your gut for probability. I'm sure the professor can do that just fine and be reasonably correct since he's done this for so long, but we've only started doing this, and I don't feel like my "gut" is developed enough to be able to approximate probabilities well.

I agree, I'm not sure where the 1/2 came from and I think saying 50% is more intuitive for most people than a probability of 1/2

this is interesting and useful

This section sums up a nice confidence interval for how reasonable we think our estimates are, but I don't see how this is going to improve our approximations

I'm still not as comfortable with the "odds" notation as I am with normal fractions or percentages to represent probabilities so I like that you wrote it both ways here.

I find this odds ratio a bit harder to use intuitively than something like 1:1 or 9:1.

wait, I am confused as to where this 2/3 came from

I think we basically just decided that it felt about right.

and information (or evidence)

$I \equiv$ What I know about the area *before* using divide and conquer.

Now I want to find the conditional probability $P(H|I)$ – namely, the probability of H given my knowledge before trying divide and conquer. There is no known algorithm for computing a probability in such a complicated problem situation. How, for example, does one represent my state of knowledge? In these cases, the best we can do is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my intuitive knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300 \ldots 3000\,\text{cm}^2$. The surprise suggests that $P(H|I)$ is larger than 1/2. The mildness of the surprise suggests that $P(H|I)$ is not much larger than 1/2. I'll quantify it as $P(H|I) = 2/3$: I would give 2-to-1 odds that the true area is within the plausible range.

Furthermore, I'll use this probability or odds to define a plausible range: It is the range for which I think 2-to-1 odds is fair. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will later help give plausible ranges an intuitive interpretation (as a region on a log-normal distribution). That interpretation will then help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly 2/3 of the probability. With a further assumption of symmetry, the plausible range $300 \ldots 3000\,\text{cm}^2$ represents the following probabilities:

$$P(A < 300\,\text{cm}^2) = 1/6;$$
$$P(300\,\text{cm} \le A \le 3000\,\text{cm}^2) = 2/3;$$
$$P(A > 3000\,\text{cm}^2) = 1/6.$$

Here is the corresponding picture with width proportional to probability:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
A	$< 300\,\text{cm}^2$	$300 \ldots 3000\,\text{cm}^2$	$> 3000\,\text{cm}^2$

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are

what do you mean about these odd's? 2-1 vs 1-1?

yeah, like I said earlier, i think speaking in terms of percentages would make it much easier to follow the probabilities that you are talking about.

Why?

This paragraph is a bit dense. Maybe you could just say that Choosing 2-1 instead of another odds ratio leads to easier interpretation...

I agree. I think that breaking down the second sentence would help a lot, it's currently kind of long and a bit haphazard

I agree with the other comment about hand-waving. While you know it will help later, if we were working on this problem, we wouldn't...

This seems a bit like hand waving to me.

i agree. it makes sense but i'm not so easily convinced to subscribe to this. if it's going to be a chapter, i want to be convinced more.

Yeah, that is true. Perhaps some more on this topic owuld be nice.

I don't entirely understand this section but it seems weirdly analytical for an estimate based on your gut...

Yeah, this section makes it seem like you knew that 2 to 1 odds is useful for plotting purposes (and has the name 'plausible range'), and you therefore found a range that fit this certainty, rather than finding a range and fitting a certainty as the prior section suggests. It's all just jumbled enough to have forced me to read it a couple of times.

I fell like this paragraph takes a lot away from your explanation...you're all but saying it really doesn't matter what's right, so i'm going to use this. it's frustrating.

assuming its equally likely to be bigger or smaller. or is that inherent by defining the range in the center of your likelihood distribution?

I really like this example and it is helping me understand the concept pretty well.

and information (or evidence)

$I \equiv$ What I know about the area *before* using divide and conquer.

Now I want to find the conditional probability $P(H|I)$ – namely, the probability of H given my knowledge before trying divide and conquer. There is no known algorithm for computing a probability in such a complicated problem situation. How, for example, does one represent my state of knowledge? In these cases, the best we can do is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my intuitive knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300 \ldots 3000\,\text{cm}^2$. The surprise suggests that $P(H|I)$ is larger than 1/2. The mildness of the surprise suggests that $P(H|I)$ is not much larger than 1/2. I'll quantify it as $P(H|I) = 2/3$: I would give 2-to-1 odds that the true area is within the plausible range.

Furthermore, I'll use this probability or odds to define a plausible range: It is the range for which I think 2-to-1 odds is fair. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will later help give plausible ranges an intuitive interpretation (as a region on a log-normal distribution). That interpretation will then help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly 2/3 of the probability. With a further assumption of symmetry, the plausible range $300 \ldots 3000\,\text{cm}^2$ represents the following probabilities:

$$P(A < 300\,\text{cm}^2) = 1/6;$$
$$P(300\,\text{cm} \le A \le 3000\,\text{cm}^2) = 2/3;$$
$$P(A > 3000\,\text{cm}^2) = 1/6.$$

Here is the corresponding picture with width proportional to probability:

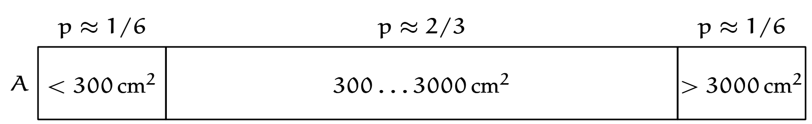

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are

This method for decision making makes a lot of sense.

I agree, I think the whole book would flow very nicely if there were many charts and tables like this.

I agree - the preceding paragraphs were somewhat confusing, but this set of equations and the chart below helped a lot.

I agree - the preceding paragraphs were somewhat confusing, but this set of equations and the chart below helped a lot.

i think the picture is much easier to comprehend than this table and the picture could stand alone.

Okay this is notation much more familiar to me from taking probability. So I guess without even knowing it, we're using continuous random variables. Which makes sense for this problem.

Well the reason we're using continuous random variables and not discrete random variables is because we're estimating the Area to be within a certain range, and not assigning it a perfect value. In the course we've always had our approximations within a range of values, so using continuous random variables is the clear choice.

I really like that this is here... It might seem kind of obvious to some but I think it's useful in defining our parameters.

Agreed. This serves as a good illustration of gut-thinking.

and information (or evidence)

$I \equiv$ What I know about the area *before* using divide and conquer.

Now I want to find the conditional probability $P(H|I)$ – namely, the probability of H given my knowledge before trying divide and conquer. There is no known algorithm for computing a probability in such a complicated problem situation. How, for example, does one represent my state of knowledge? In these cases, the best we can do is to introspect or, in plain English, to talk to our gut.

My gut is the organ with the most access to my intuitive knowledge and its incompleteness, and it tells me that I would feel mild surprise but not shock if I learned that the true area lay outside the range $300 \ldots 3000\,\text{cm}^2$. The surprise suggests that $P(H|I)$ is larger than 1/2. The mildness of the surprise suggests that $P(H|I)$ is not much larger than 1/2. I'll quantify it as $P(H|I) = 2/3$: I would give 2-to-1 odds that the true area is within the plausible range.

Furthermore, I'll use this probability or odds to define a plausible range: It is the range for which I think 2-to-1 odds is fair. I could have used a 1-to-1 odds range instead, but the 2-to-1 odds range will later help give plausible ranges an intuitive interpretation (as a region on a log-normal distribution). That interpretation will then help quantify how to combine plausible ranges.

For the moment, I need only the idea that the plausible range contains roughly 2/3 of the probability. With a further assumption of symmetry, the plausible range $300 \ldots 3000\,\text{cm}^2$ represents the following probabilities:

$$P(A < 300\,\text{cm}^2) = 1/6;$$
$$P(300\,\text{cm} \le A \le 3000\,\text{cm}^2) = 2/3;$$
$$P(A > 3000\,\text{cm}^2) = 1/6.$$

Here is the corresponding picture with width proportional to probability:

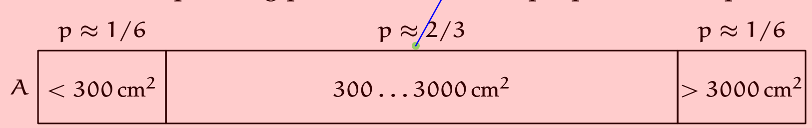

For the height h and width w, after doing divide and conquer and using the similarity between A4 and letter paper, the plausible ranges are

This is really useful to those of us who haven't had much probability, I think some diagrams like this in the previous section would be helpful to its understanding

Although the jumping between "odds" and "probability" is kind of confusing. I think I'd prefer it if you stuck with just one (probability, preferably)

The "jump" between odds and probability really shouldn't be a jump at all. Here's a quick explanation that should simplify things:

Let p the probability that an event occurs. Then the probability that the event does not occur = 1 - p. The odds of an event occurring refers to the ratio of the probability that the event occurs divided by the probability that the event does not occur. That is, the odds = p/(1-p).

This is a good explanation from 8:55. Thanks.

I agree. Odds are simply a more everyday way of talking about a probability. They are directly related to each other so it shouldn't be too much to switch between the two.

I think that using both odds & probability is a good thing...I totally get the probability better, but odds are used more in everyday life and we need to get a better understanding of them anyway

I agree. If such diagrams were available in the previous section, I think I might have absorbed much more.

I'd expect a picture to be more like a graph or histogram, but this is useful!

28 . . . 32 cm and 19 . . . 21 cm respectively. Here are their probability interpretations:

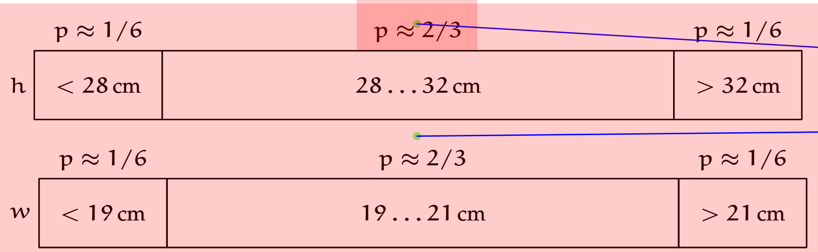

Computing the plausible range for the area requires a complete probabilistic description of a plausible range. There is an answer to this question that depends on the information available to the person giving the range. But no one knows the exact recipe to deduce probabilities from the complex, diffuse, seemingly contradictory information lodged in a human mind.

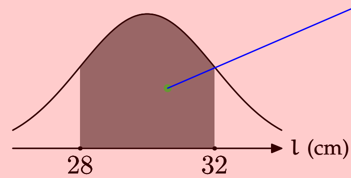

The best that we can do for now is to guess a reasonable and convenient probability distribution. I will use a log-normal distribution, meaning that the uncertainty in the quantity's logarithm has a normal (or Gaussian) distribution. As an example, the figure shows the probability distribution for the length of A4 length (after taking into account the similarity to letter paper). The shaded range is the the so-called one-sigma range $\mu-\sigma$ to $\mu+\sigma$. It contains 68% of the probability – a figure conveniently close to 2/3. So to convert a plausible range to a log-normal distribution, use the lower and upper endpoints of the plausible range as $\mu - \sigma$ to $\mu + \sigma$. The peak of the distribution – the most likely value – occurs midway between the endpoints. Since 'midway' is on a logarithmic scale, the midpoint is at $\sqrt{28 \times 32}$ cm or approximately 29.93 cm.

The log-normal distribution supplies the missing information required to combine plausible ranges. When adding independent quantities, you add their means and their variances. So when multiplying independent quantities, add the means and variances in the logarithmic space.

Here is the resulting recipe. Let the plausible range for h be $l_1 \ldots u_1$ and the plausible range for w be $l_2 \ldots u_2$. First compute the geometric mean (midpoint) of each range:

So every time we guess a range, we use the plausible range probability distribution (1/6, 2/3, 1/6), because all of our guts are probably 2/3 right?

why is the probability the same?

I like these charts, very clear!

I agree - it really makes the point about how the probabilities interrelate.

This wording of this sentence is kind of confusing to me.

This is confusing to me - I'm not sure why we have to relate talking to our gut to probability if it is really impossible to come up with a good probability

Yet another nice diagram!!

The diagrams from this reading have been very helpful.

28 . . . 32 cm and 19 . . . 21 cm respectively. Here are their probability interpretations:

	p ≈ 1/6	p ≈ 2/3	p ≈ 1/6
h	< 28 cm	28 . . . 32 cm	> 32 cm

	p ≈ 1/6	p ≈ 2/3	p ≈ 1/6
w	< 19 cm	19 . . . 21 cm	> 21 cm

Computing the plausible range for the area requires a complete probabilistic description of a plausible range. There is an answer to this question that depends on the information available to the person giving the range. But no one knows the exact recipe to deduce probabilities from the complex, diffuse, seemingly contradictory information lodged in a human mind.

The best that we can do for now is to guess a reasonable and convenient probability distribution. I will use a log-normal distribution, meaning that the uncertainty in the quantity's logarithm has a normal (or Gaussian) distribution. As an example, the figure shows the probability distribution for the length of A4 length (after taking into account the similarity to letter paper). The shaded range is the the so-called one-sigma range $\mu-\sigma$ to $\mu+\sigma$. It contains 68% of the probability – a figure conveniently close to 2/3. So to convert a plausible range to a log-normal distribution, use the lower and upper endpoints of the plausible range as $\mu-\sigma$ to $\mu+\sigma$. The peak of the distribution – the most likely value – occurs midway between the endpoints. Since 'midway' is on a logarithmic scale, the midpoint is at $\sqrt{28\times 32}$ cm or approximately 29.93 cm.

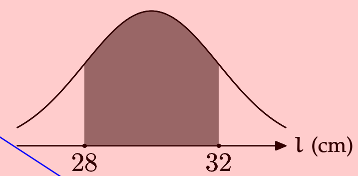

The log-normal distribution supplies the missing information required to combine plausible ranges. When adding independent quantities, you add their means and their variances. So when multiplying independent quantities, add the means and variances in the logarithmic space.

Here is the resulting recipe. Let the plausible range for h be $l_1 \ldots u_1$ and the plausible range for w be $l_2 \ldots u_2$. First compute the geometric mean (midpoint) of each range:

why would we use this over just a normal one?

It seems like this method keeps the error from doubling when we multiply the length and the area? Log-normal graphs are used a lot in probability.

also, since each length is clearly bounded below by 0, the log-normal is more appropriate.

A log normal also centers the range about the geometric mean, not arithmetic mean, which is much more appropriate for this type of problem.

It also doesn't look particularly different from a normal distribution when the range is concentrated away from 0.

Those are indeed the reasons. To really see the difference between normal and log normal, I should use a range like 10^2 .. 10^4 (i.e. 10^3 give or take a factor of 10).

What we mean intuitively is that 1000 is the most likely value, and 300 is about as likely as 3333 (=10000/3). That is the log-normal interpretation.

In the normal interpretation, the most likely value would be 5050, which is already very strange. And there would be a significant probability of the value being less than 0 (because 0 is only about 1 sigma below the mean of 5050).

It would be nice if log-normal was defined at some point, for those of us who may not have encountered the term before.

Nevermind. It is... I just assumed because that was the second time I'd seen it. Should have read further.

No, I agree – it's not really that well defined even later on the page.

This seems a little far to go just to estimate the area of a piece of paper

True, but it will be really useful in real life estimations, if we don't know values and it becomes important to estimate them exactly

It's a simple example of a two-parameter problem.

I'd still rather examine an example where I wouldn't just guess length and width and error.

28 . . . 32 cm and 19 . . . 21 cm respectively. Here are their probability interpretations:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
h	< 28 cm	28 . . . 32 cm	> 32 cm

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
w	< 19 cm	19 . . . 21 cm	> 21 cm

Computing the plausible range for the area requires a complete probabilistic description of a plausible range. There is an answer to this question that depends on the information available to the person giving the range. But no one knows the exact recipe to deduce probabilities from the complex, diffuse, seemingly contradictory information lodged in a human mind.

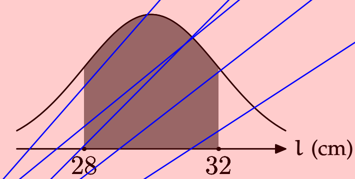

The best that we can do for now is to guess a reasonable and convenient probability distribution. I will use a log-normal distribution, meaning that the uncertainty in the quantity's logarithm has a normal (or Gaussian) distribution. As an example, the figure shows the probability distribution for the length of A4 length (after taking into account the similarity to letter paper). The shaded range is the the so-called one-sigma range $\mu - \sigma$ to $\mu + \sigma$. It contains 68% of the probability – a figure conveniently close to 2/3. So to convert a plausible range to a log-normal distribution, use the lower and upper endpoints of the plausible range as $\mu - \sigma$ to $\mu + \sigma$. The peak of the distribution – the most likely value – occurs midway between the endpoints. Since 'midway' is on a logarithmic scale, the midpoint is at $\sqrt{28 \times 32}$ cm or approximately 29.93 cm.

The log-normal distribution supplies the missing information required to combine plausible ranges. When adding independent quantities, you add their means and their variances. So when multiplying independent quantities, add the means and variances in the logarithmic space.

Here is the resulting recipe. Let the plausible range for h be $l_1 \ldots u_1$ and the plausible range for w be $l_2 \ldots u_2$. First compute the geometric mean (midpoint) of each range:

I wish my high school stats class had taught me cool cases like the ones you use as an example.

This is a very interesting idea and easily applicable to many real life examples.

I agree, the analysis here is very clear and well described.

you should just mention that mu is the mean before putting this out there

It might be helpful if you stated that mu was the mean and sigma the standard deviation.

I think that you should introduce the 68-95-99 rule here.

Ah its all starting to come together. A standard deviation away from the mean on both sides is 2/3 which is what he set to be the probability of our range. So I guess in general, that's what we should always set the probability to.

While I understand his motivation for setting the probability to 2/3 in this case since he wanted to give 2-to-1 odds and it also was very close to the 68% confidence interval encompassed by being one standard deviation (or alternatively, one z-score) away from the mean, he also did it because 2-to-1 odds was the appropriate probability in this problem. If the odds really aren't 2-to-1, it would be unwise to "always set the probability to 2/3" as you mentioned.

what do you mean it is on a log scale?

I think it's also fair to add that sometimes you will guess too low and sometimes you will guess to high. This, over time, will even out.

I feel like using a log scale is unnecessary and complicates things.

I was also unsure of why we use a log-normal scale. I understand the need for a normal / Gaussian graph, but why a logarithmic one? Is there something in this problem that helps us identify when to use which?

Why not just use the average of your two ranges and multiply those to get your guess?

Is this an approximation or a definition of the midpoint of a log graph? I'm not really up on statistics.

28 . . . 32 cm and 19 . . . 21 cm respectively. Here are their probability interpretations:

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
h	< 28 cm	28 . . . 32 cm	> 32 cm

	$p \approx 1/6$	$p \approx 2/3$	$p \approx 1/6$
w	< 19 cm	19 . . . 21 cm	> 21 cm

Computing the plausible range for the area requires a complete probabilistic description of a plausible range. There is an answer to this question that depends on the information available to the person giving the range. But no one knows the exact recipe to deduce probabilities from the complex, diffuse, seemingly contradictory information lodged in a human mind.

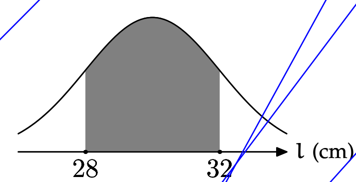

The best that we can do for now is to guess a reasonable and convenient probability distribution. I will use a log-normal distribution, meaning that the uncertainty in the quantity's logarithm has a normal (or Gaussian) distribution. As an example, the figure shows the probability distribution for the length of A4 length (after taking into account the similarity to letter paper). The shaded range is the the so-called one-sigma range $\mu - \sigma$ to $\mu + \sigma$. It contains 68% of the probability – a figure conveniently close to 2/3. So to convert a plausible range to a log-normal distribution, use the lower and upper endpoints of the plausible range as $\mu - \sigma$ to $\mu + \sigma$. The peak of the distribution – the most likely value – occurs midway between the endpoints. Since 'midway' is on a logarithmic scale, the midpoint is at $\sqrt{28 \times 32}$ cm or approximately 29.93 cm.

The log-normal distribution supplies the missing information required to combine plausible ranges. When adding independent quantities, you add their means and their variances. So when multiplying independent quantities, add the means and variances in the logarithmic space.

Here is the resulting recipe. Let the plausible range for h be $l_1 \ldots u_1$ and the plausible range for w be $l_2 \ldots u_2$. First compute the geometric mean (midpoint) of each range:

I understand why you would use this number in approximating, but doesn't this method complicate the overall estimating process? Also, why couldn't we have just used (28+32)/2 = 30cm ?

I'm in 18.440 and we just went over this concept, it's cool to see it applied here!

So I was following the reading until now.... Any 18.440 students wish to explain this further?

What exactly do you mean by variances here?

I'm glad we finally got to this point. It means that the level of uncertainty continues to increase as we continue to multiply quantities, assuming that there is no new information into the system.

I hope we go over this in class

I'm confused as to why h and w are described in both l and u? Why isn't h described as l_1... l_2 and w as u_1... u_2? The combination of variables is confusing and unintuitive to me.

l and u refer to lower and upper bounds, I assume.

I agree, the use of variables here is a bit strange

$$\mu_1 = \sqrt{l_1 u_1};$$
$$\mu_2 = \sqrt{l_2 u_2}.$$

The midpoint of the range for $A = hw$ is the product of the two midpoints:

$$\mu = \mu_1 \mu_2. \qquad (8.9)$$

To compute the plausible range, first compute the ratios measuring the width of the ranges:

$$r_1 = u_1/l_1;$$
$$r_2 = u_2/l_2.$$

These ratios measure the width of the ranges. The combined ratio – that is, the ratio of endpoints for the combined plausible range – is

$$r = \exp\left(\sqrt{(\ln r_1)^2 + (\ln r_2)^2}\right).$$

For approximate range calculations, the following contour graphs often provide enough accuracy:

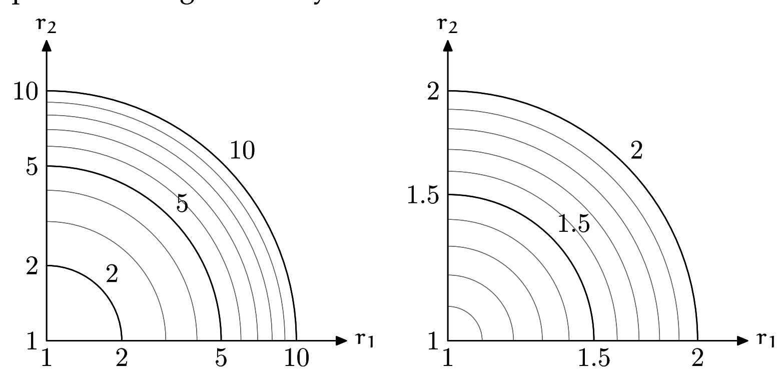

After finding the range, choose the lower and upper endpoints l and u to make $u/l = r$ and $\sqrt{lu} = \mu$. In other words, the plausible range is

$$\frac{\mu}{\sqrt{r}} \ldots \mu\sqrt{r}.$$

Problem 8.1 Deriving the ratio
Use Bayes theorem to confirm this method for combining plausible ranges.

Let's check this method in a simple example where the width and height ranges are $1 \ldots 2\,\text{m}$. What is the plausible range for the area? The naive

Comments on page 5

I"m pretty lost as well. Also not sure why we choose to use two approximations and then combine them.

I got very lost here...and I still am. I get what you're saying on the next page, but (even after sleeping on it) I just can't seem to follow what you're doing here.

I got kind of lost in here. I think it might be the l, u notations. They don't seem very intuitive.

I think l stands for lower and u for upper. But I had to read it a few times before I saw that.

Good call! thanks that really helped me make sense of it when I was reading it

While that does make sense, the mu used for midpoint is very similar to the u. I think it would be less confusing if the u and l were capitalized.

you lost me from here down...

I am confused on what you are talking about

Could this be explained graphically as well?

$$\mu_1 = \sqrt{l_1 u_1};$$
$$\mu_2 = \sqrt{l_2 u_2}.$$

The midpoint of the range for $A = hw$ is the product of the two midpoints:

$$\mu = \mu_1 \mu_2. \qquad (8.9)$$

To compute the plausible range, first compute the ratios measuring the width of the ranges:

$$r_1 = u_1/l_1;$$
$$r_2 = u_2/l_2.$$

These ratios measure the width of the ranges. The combined ratio – that is, the ratio of endpoints for the combined plausible range – is

$$r = \exp\left(\sqrt{(\ln r_1)^2 + (\ln r_2)^2}\right).$$

For approximate range calculations, the following contour graphs often provide enough accuracy:

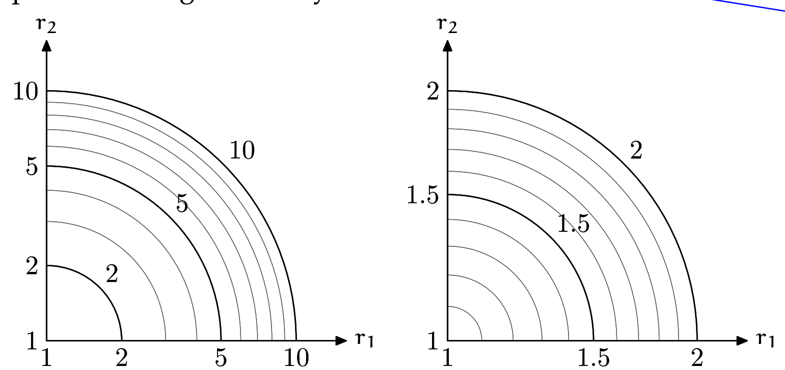

After finding the range, choose the lower and upper endpoints l and u to make $u/l = r$ and $\sqrt{lu} = \mu$. In other words, the plausible range is

$$\frac{\mu}{\sqrt{r}} \ldots \mu\sqrt{r}.$$

Problem 8.1 Deriving the ratio
Use Bayes theorem to confirm this method for combining plausible ranges.

Let's check this method in a simple example where the width and height ranges are $1 \ldots 2\,\text{m}$. What is the plausible range for the area? The naive

This is the most confusing part. I don't know where this came from.

Yeah I feel like all of this suddenly comes out of nowhere - more exposition would have been nice.

i agree.

Yes, agreed. A bit more explanation for those of us who haven't taken a probability class would be great.

I would add that an overview on log's as they relate to this class would help me with a lot of this stuff. I learned the change of base formula and basic manipulations back in freshmen year, but it seems to me we use them a bit more extensively in this class.

This follows from the fact that variances of the logs add for log-normal distributions.

If we had assumed these were normal distributions, we would simply have s=sqrt(s1^2 + s2^2), where s=standard deviation instead of endpoints ratio.

Unfortunately, I feel like most of this came out of nowhere. Like why we can add the means and variances on the log space or something.

So, on a log-normal plot, the variance is the square of the log of the standard deviation of the original data? I think that makes sense, but it would probably be worth explaining.

I'm not familiar with these plots...what are they showing/proving exactly? Does the normal distribution above not suffice?

Yeah...what is this? More explanation is needed.

I think it's just an easy way to find r from r1 and r2. The magnitude of the contour at (r1,r2) is r.

I'm also confused as to what these are showing. More explanation would probably be helpful.

I haven't really encountered these a lot... a little more explanation might be nice?

I think that this problem just has a few really complicated concepts. The first was about multiplying independent quantities, and now the contour graphs and the ratio analysis.

$$\mu_1 = \sqrt{l_1 u_1};$$
$$\mu_2 = \sqrt{l_2 u_2}.$$

The midpoint of the range for $A = hw$ is the product of the two midpoints:

$$\mu = \mu_1 \mu_2. \tag{8.9}$$

To compute the plausible range, first compute the ratios measuring the width of the ranges:

$$r_1 = u_1/l_1;$$
$$r_2 = u_2/l_2.$$

These ratios measure the width of the ranges. The combined ratio – that is, the ratio of endpoints for the combined plausible range – is

$$r = \exp\left(\sqrt{(\ln r_1)^2 + (\ln r_2)^2}\right).$$

For approximate range calculations, the following contour graphs often provide enough accuracy:

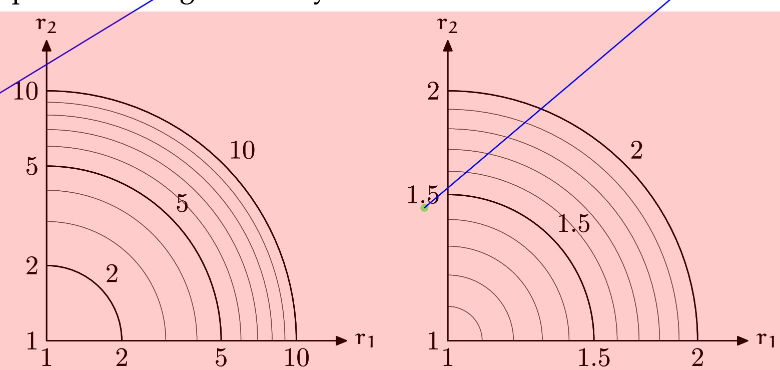

After finding the range, choose the lower and upper endpoints l and u to make $u/l = r$ and $\sqrt{lu} = \mu$. In other words, the plausible range is

$$\frac{\mu}{\sqrt{r}} \ldots \mu\sqrt{r}.$$

Problem 8.1 Deriving the ratio
Use Bayes theorem to confirm this method for combining plausible ranges.

Let's check this method in a simple example where the width and height ranges are $1 \ldots 2\,\text{m}$. What is the plausible range for the area? The naive

I'm pretty confused, as I think others are too...this seems like a lot of work to do – when should we actually use this method? I'm also afraid I would not be able to implement it correctly.

agreed. i also don't know what to make of these plots. kind of confusing.

I agree. I don't really understand what these graphs are supposed to be used for.

Is this necessary just to prove why divide and conquer works??

A lot of people, including myself, seem to struggle to understand what exactly the meaning of these graphs are. I'm hoping he'll explain it further in class.

Yeah, I'm also confused here. Are these graphs necessary? What exactly do they represent, and what are you trying to show with them?

These graphs are not necessary, but I personally actually like them. They are just a demonstration of how ot find r from r1 and r2 within a reasonable accuracy.

I think this is helpful for the more theoretically oriented people who want an explanation. As an engineer, I don't care how things work, as long as they work. Why worry about reinventing the wheel. I wouldn't bother reading this chapter if I bought this book, but I think it's necessary for a class to include this explanation

$$\mu_1 = \sqrt{l_1 u_1};$$
$$\mu_2 = \sqrt{l_2 u_2}.$$

The midpoint of the range for $A = hw$ is the product of the two midpoints:

$$\mu = \mu_1 \mu_2. \quad (8.9)$$

To compute the plausible range, first compute the ratios measuring the width of the ranges:

$$r_1 = u_1/l_1;$$
$$r_2 = u_2/l_2.$$

These ratios measure the width of the ranges. The combined ratio – that is, the ratio of endpoints for the combined plausible range – is

$$r = \exp\left(\sqrt{(\ln r_1)^2 + (\ln r_2)^2}\right).$$

For approximate range calculations, the following contour graphs often provide enough accuracy:

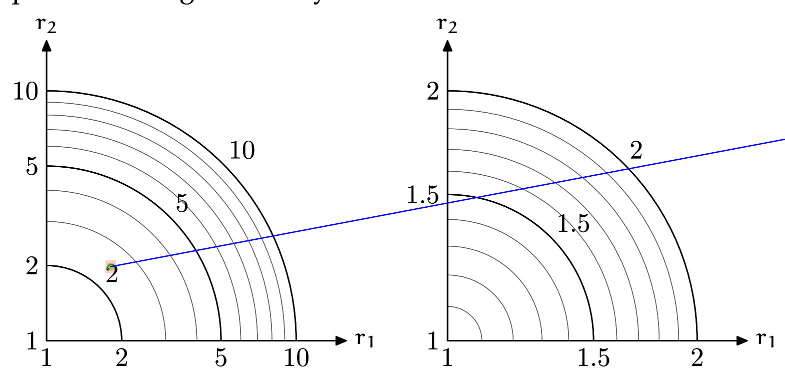

After finding the range, choose the lower and upper endpoints l and u to make $u/l = r$ and $\sqrt{lu} = \mu$. In other words, the plausible range is

$$\frac{\mu}{\sqrt{r}} \ldots \mu\sqrt{r}.$$

Problem 8.1 Deriving the ratio
Use Bayes theorem to confirm this method for combining plausible ranges.

Let's check this method in a simple example where the width and height ranges are 1 . . . 2 m. What is the plausible range for the area? The naive

What is the difference between these two graphs? (I see the numbers are different, but I"m really confused as to what this shows)

Yeah, I'm pretty lost here as well–what do these two show? the upper and lower ranges?

I also don't quite get what point the graphs are supposed to show.

I think they are illustrations for two different ranges so you can see what a log plot looks like in this contour form. The second plot is the lower left quarter-circle enlarged. I think you take the two endpoints of the range, then it would give you a sense of the likeliest value. So if r1 = 2 and r2 = 5, that point on the graph is somewhere just outside the 5 contour line, so r 5.5? I am still pretty confused about the use of these graphs though.

I'm confused about these graphs as well.

that explanation, that one a blow-up of the other helps, but it still doesn't really make the point very clear.

Clearly I need to explain this better. I'll have a go in lecture and hopefully that will result in a better explanation for the book too.

So for the problem below, the range is somewhere like here? I guess that looks like 2.67

What do the varying contours mean? Don't really understand the significance of these diagrams..

is this really a process you would undertake to simply estimate the area of a sheet of paper?

It's definitely overkill for the sheet of paper. But the process is general, and if you want to know how accurate an estimate is (i.e. how confident you can be about an estimate), now you know how to do it. In lecture we'll do a more substantial example, and maybe I should also put that in the reading.

I got confused here that we were no longer still talking about the paper since these numbers seemed too big for that. But after reading a bit more, I see that this was just an example of easy cases to test the theory. You could mention that to tie this back more clearly to earlier sections

approach of simply multiplying endpoints produces a plausible range of $1 \ldots 4\,\text{m}^2$ – a width of a factor of 4. However, this range is too pessimistic and the correct range should be narrower. Using the log-normal distribution, the range spans a factor of

$$\exp\left(\sqrt{2 \times (\ln 2)^2}\right) \approx 2.67.$$

This span and the midpoint determine the range. The area midpoint is the product of the width and height midpoints, each of which is $\sqrt{2}\,\text{m}$. So the midpoint is $2\,\text{m}^2$. The correct endpoints of the plausible range are therefore

$$\frac{2\,\text{m}^2}{\sqrt{2.67}} \ldots 2\,\text{m}^2 \times \sqrt{2.67}$$

or $1.23 \ldots 3.27\,\text{m}^2$. In other words, I assign roughly a 1/6 probability that the area is less than $1.23\,\text{m}^2$ and roughly a 1/6 probability that it is greater than $3.27\,\text{m}^2$. Those conclusions seem reasonable when using such uncertain knowledge of length and width.

Having checked that the method is reasonable, it is time to test it in the original illustrative problem: the plausible area range for an A4 sheet. The naive plausible range was $532 \ldots 672\,\text{cm}^2$, and the correct plausible range will be narrower. Indeed, the log-normal method gives the narrower area range of $550 \ldots 650\,\text{cm}^2$ with a best guess (most likely value) of $598\,\text{cm}^2$. How did we do? The true area is exactly $2^{-4}\,\text{m}^2$ or $625\,\text{cm}^2$ because – I remembered only after doing this calculation! – An paper is constructed to have one-half the area of A$(n-1)$ paper, with A0 paper having an area of $1\,\text{m}^2$. The true area is only 5% larger than the best guess, suggesting that we used accurate information about the length and width; and it falls within the plausible range but not right at the center, suggesting that the method for computing the plausible range is neither too daring nor too conservative.

The analysis of combining ranges illustrates the two crucial points about divide-and-conquer reasoning. First, the main benefit comes from subdividing vague knowledge (such as the area itself) into pieces about which our knowledge is accurate (the length and the width). Second, this benefit swamps the small penalty in accuracy that results from combining many quantities together.

Thanks for giving us another example!

Why is the extra square root of 2.67 taken here? I thought the span was just 2.67 ˆ 1.

I was a little confused at times about when we stopped talking about this, and it seemed to take a while to get back to this, although it brings the points together very clearly when it is brought up again, I think that it could be sooner/ more integrated into the preceding examples.

Is there an example that's just as simple and illustrative but maybe a little more relevant or ..interesting than the area of a sheet of European paper? (no offense to European paper or anything)

I am really confused on why we are doing this probability calculations? what are they actually telling us?

I think this reading overall was very helpful and informative. I do always get a little intimidated when a lot of equations and numbers are thrown in.

Given that this represents a change of about 3% from the original endpoints, it seems sort of silly to go through all the log-normal calculations. Wouldn't the original range be "good enough" for most practices?

This notation was confusing at first, maybe because I'm not used to referring to paper size like this. I'm not sure how to make it clearer though, maybe with an example?

I was also confused as to why you were saying "An paper" ... breaking up this sentence would probably help a lot...or include a chart with A0, A1, A2, ...

Very interesting fact. explains why each size is double the previous, although i had no idea A0 was 1mˆ2.

Huh...really? I wasn't even aware that there was A2 or A3 paper. That's pretty cool.

Yeah, this seems to make a bit more sense to me now.

This is a really cool piece of semi useful knowledge.

there is so much rhyme and reason to metric.

approach of simply multiplying endpoints produces a plausible range of $1 \ldots 4\,\mathrm{m}^2$ – a width of a factor of 4. However, this range is too pessimistic and the correct range should be narrower. Using the log-normal distribution, the range spans a factor of

$$\exp\left(\sqrt{2 \times (\ln 2)^2}\right) \approx 2.67.$$

This span and the midpoint determine the range. The area midpoint is the product of the width and height midpoints, each of which is $\sqrt{2}\,\mathrm{m}$. So the midpoint is $2\,\mathrm{m}^2$. The correct endpoints of the plausible range are therefore

$$\frac{2\,\mathrm{m}^2}{\sqrt{2.67}} \ldots 2\,\mathrm{m}^2 \times \sqrt{2.67}$$

or $1.23 \ldots 3.27\,\mathrm{m}^2$. In other words, I assign roughly a 1/6 probability that the area is less than $1.23\,\mathrm{m}^2$ and roughly a 1/6 probability that it is greater than $3.27\,\mathrm{m}^2$. Those conclusions seem reasonable when using such uncertain knowledge of length and width.

Having checked that the method is reasonable, it is time to test it in the original illustrative problem: the plausible area range for an A4 sheet. The naive plausible range was $532 \ldots 672\,\mathrm{cm}^2$, and the correct plausible range will be narrower. Indeed, the log-normal method gives the narrower area range of $550 \ldots 650\,\mathrm{cm}^2$ with a best guess (most likely value) of $598\,\mathrm{cm}^2$. How did we do? The true area is exactly $2^{-4}\,\mathrm{m}^2$ or $625\,\mathrm{cm}^2$ because – I remembered only after doing this calculation! – An paper is constructed to have one-half the area of A$(n-1)$ paper, with A0 paper having an area of $1\,\mathrm{m}^2$. The true area is only 5% larger than the best guess, suggesting that we used accurate information about the length and width; and it falls within the plausible range but not right at the center, suggesting that the method for computing the plausible range is neither too daring nor too conservative.

The analysis of combining ranges illustrates the two crucial points about divide-and-conquer reasoning. First, the main benefit comes from subdividing vague knowledge (such as the area itself) into pieces about which our knowledge is accurate (the length and the width). Second, this benefit swamps the small penalty in accuracy that results from combining many quantities together.

I thought this method was overkill, applied here to estimating paper area, but it served as an easy to follow example of how to apply it, even if the example itself wasn't necessarily the most interesting one to use it on.

I think this was the point of doing it in 'slow motion' hopefully we will also see it applied to more exciting examples later in the book.

It was nice to have an easy example, but at certain points the "slow motion" ran a little too slow so I lost site of the motivation for much of the example. As the original commenter suggested, maybe a slightly more relatable example would have been helpful.

It seems like we've been doing this without realizing it for a while, is the extra time needed with this method ever worth it over the speed of trusting one's gut? (assuming of course we have a reasonable gut intuition)

Well, all semester people have been complaining that they are not sure of their answers and their guts are giving them really bit ranges...so here is a way to narrow it down and explain exactly what you should do when asked to make a gut reasoning.

I was wondering the same. This seems like an excessive amount of work for an approximation.

I agree. This is making some of the approximations that we did at the beginning of the term make mathematical sense.

I think that most of the work in this section has been to look at the errors in what you're estimating...these are not actually calculations that you would be doing, unless you need to know how far off your approximations are.

I really like this conclusion. Great summary!

I really like this conclusion. Great summary!

Im not really sure i got this from the reading...

approach of simply multiplying endpoints produces a plausible range of $1\ldots4\,\mathrm{m}^2$ – a width of a factor of 4. However, this range is too pessimistic and the correct range should be narrower. Using the log-normal distribution, the range spans a factor of

$$\exp\left(\sqrt{2\times(\ln 2)^2}\right)\approx 2.67.$$

This span and the midpoint determine the range. The area midpoint is the product of the width and height midpoints, each of which is $\sqrt{2}\,\mathrm{m}$. So the midpoint is $2\,\mathrm{m}^2$. The correct endpoints of the plausible range are therefore

$$\frac{2\,\mathrm{m}^2}{\sqrt{2.67}}\ldots 2\,\mathrm{m}^2\times\sqrt{2.67}$$

or $1.23\ldots3.27\,\mathrm{m}^2$. In other words, I assign roughly a 1/6 probability that the area is less than $1.23\,\mathrm{m}^2$ and roughly a 1/6 probability that it is greater than $3.27\,\mathrm{m}^2$. Those conclusions seem reasonable when using such uncertain knowledge of length and width.

Having checked that the method is reasonable, it is time to test it in the original illustrative problem: the plausible area range for an A4 sheet. The naive plausible range was $532\ldots672\,\mathrm{cm}^2$, and the correct plausible range will be narrower. Indeed, the log-normal method gives the narrower area range of $550\ldots650\,\mathrm{cm}^2$ with a best guess (most likely value) of $598\,\mathrm{cm}^2$. How did we do? The true area is exactly $2^{-4}\,\mathrm{m}^2$ or $625\,\mathrm{cm}^2$ because – I remembered only after doing this calculation! – An paper is constructed to have one-half the area of A$(n-1)$ paper, with A0 paper having an area of $1\,\mathrm{m}^2$. The true area is only 5% larger than the best guess, suggesting that we used accurate information about the length and width; and it falls within the plausible range but not right at the center, suggesting that the method for computing the plausible range is neither too daring nor too conservative.

The analysis of combining ranges illustrates the two crucial points about divide-and-conquer reasoning. First, the main benefit comes from subdividing vague knowledge (such as the area itself) into pieces about which our knowledge is accurate (the length and the width). Second, this benefit swamps the small penalty in accuracy that results from combining many quantities together.

I wonder if this section would be better placed near the divide and conquer section.

perhaps but the divide an conquer is supposed to be really easy, and this is not so easy to start off with. it think it makes sense here too.

I think it's better suited here. At first, I think it's best to accept divide-and-conquer as an approximation method that works without us having to see proof of it. In order to understand how it works with a proof in probability, you would need to be familiar with a little probability and I like having it in this section. Plus, having this at the beginning of the course might scare some people with no previous exposure to probability away.

I think with the probability usage, it's rightly placed here.

I think this sentence was a very clear conclusion.

This makes a lot of sense.

wasn't quite sure how this was proved in the reading above.

His original guess was 300...3000 - a very broad guess, because our knowledge is vague. By subdividing the unknown into things we know more about (lengths), we gain some inaccuracy because we're multiplying unknowns, but less inaccuracy than simply multiplying our uncertainty (because our uncertainty in one number will to some extent be canceled by our uncertainty in the other number). And this inaccuracy is much less than what we gained by going to numbers we know more about - hence why we can get to a range of 550-650 using divide and conquer (as compared to his original 300-3000). Hence the small penalty in accuracy from using many numbers is swamped by the benefit of using those many numbers. Thus divide and conquer works.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

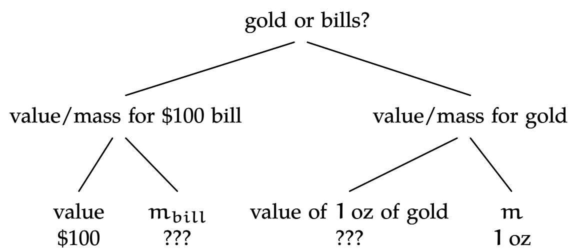

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

but the value of gold is always changing...and so is the value of the dollar!

Yeah, I remember the last value of gold to be 1100 (it's now... 1165), but I check it very often. The value of the dollar is based on gold I believe, so I think the only important thing is the variations of the gold. At this rate, the 10s and the 20s are probably not as good!

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

Having broken into a bank vault, should we take the bills or the gold?

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

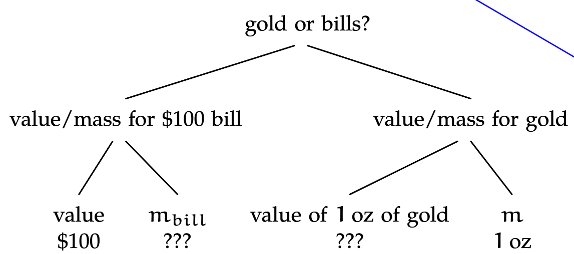

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

Read this subsection for the memo due Tuesday at 10pm. It applies the plausible-range analysis to another divide-and-conquer problem (hopefully a hypothetical one).

what about forfeited?

This is one of my favorite readings from this class so far!

Wow. I'm excited. Hope the hype holds up.

Don't like this word. Could you replace it with "forewent" or something similar?

I agree. "Forgot" sounds strange here.

I think he means "forgo"

exactly what I was thinking. forgo, not forgot!

Does this include Sanjoy?

or do you perhaps mean bank-robbing?

This is a very exciting way to introduce the section because we have all thought about it once.

I hope we do this analysis before we break into the vault

I like this problem. It's something that we can all understand. On top of that, most of us, or at least myself, don't know the answer, so it makes it more interesting to read.

there's also this sort of, i don't really know how to say it... zing, to a problem that involves a bank hest :)

makes me want to read it more... or maybe i just seen to much oceans 11.

Nah, I think we're all intrigued. I'm just waiting for the news report about MIT students getting caught robbing a bank.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

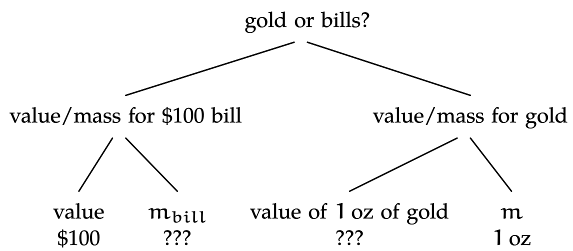

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

Without reading further, I'm guessing take the gold. Per oz, it's worth more I think.

But you need to take ease of removal into consideration, gold is much more of a hassle to remove than bills

well considering that gold is worth over $1000 an ounce and a stack of printer papers weights 20lbs (320 oz.) I would think the gold is worth more than the bills.

You also might be less likely to get caught spending the gold since you could melt it or something.

The dollar is dying, gold is getting more valuable. Also you can wash the dye that's gonna explode when you open the bag off of gold but not a dollar bill. Take the gold.

There's also that whole issue of serial numbers on the paper cash

Well, that's the problem we're facing in this approximation example. ;) Though I personally would probably take the bills, assuming they were sufficiently large.

But sufficiently large bills are often marked and can be traced. And large amounts of cash are heavy and difficult to move, too.

Although, the same could be said for gold. Perhaps it is best if we just refrain from stealing from the vault?

... or just steal electronically...

Is this going to turn into a backpack problem? Cause honestly I think that would be really cool, I haven't done anything like that in two years

this is well...off topic, but for bills, can't they put tracers on it, or that purple stuff? i suppose since it is in a vault, maybe not...but who knows.

I guess it depends on your views of dollar and gold futures too...

I guess it depends on your views of dollar and gold futures too...

do they keep gold in bank vaults? just wondering

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

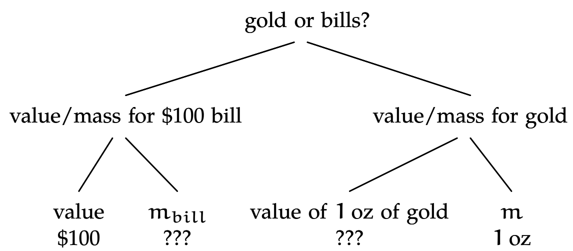

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

i don't know what this word means in this context.

Perhaps 'taking'?

I believe it means getting rid of the money in a way that is not suspicious, i.e. so you can actually use the money to buy things instead of handing out stolen $100s everywhere.

You are correct: it simply refers to selling stolen goods.

In this context...getting rid of $100 is not hard. $100,000 bills would be harder. And how would you explain giant blocks of gold?

You can always melt gold! ;)

and there are people who would take it off your hands you deal in the business so a lesser then retail price - you would still make a killing

Selling, yes.

Just a British way of saying it!

I think it's common to all dialects of English (including the British and American ones).

Is that British for laundering?

Laundering is typically money, but fencing is for non-currency goods.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

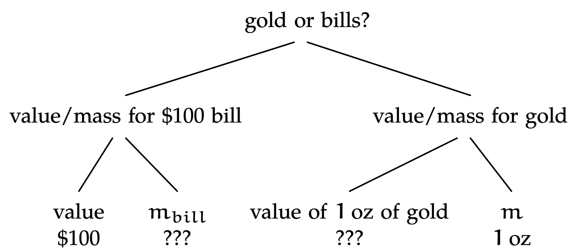

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

This is what I was thinking about, although I don't the analysis would be too difficult, just a little more estimation in terms of space and weight

mmm, I think he is referring to things such as "washing" the money/gold, so that you would actually be able to sell the money/gold without being traced/caught.

For example, taking the money/gold and using it to acquire other goods (such as stamps - something constant in value and far less suspicious). This is often difficult or costs a lot to have someone else do it for you.

Couldn't you lump this cost? I feel like we could simply apply some "washing" fraction/percentage that will account for this.

no, it means which is easier to pawn, melt, purchase with, etc. or who you know and what they deal in. it doesn't have to do with weight/space at this point. and i agree, there's no way to estimate it from a pedagogical standpoint.

Probably beyond the scope of any real textbook!

Based on weight? I would guess the bills

Hmm with my gut instinct I'm really not sure...could they actually be quite comparable?

It depends on the price gold is going for, and what methods we have of extracting the gold or bills

gold keeps on increasing in value and will likely continue to do so while the american dollar has been decreasing in value in the global economy. it really is smarter to invest in gold.

while it is smart to invest in gold and that its value keeps on increasing, i don't think that you would just rob a bank to turn around and invest your money for a period of time (i.e. put it in another bank). You probably are gonna want to spend at least a large amount of it relatively soon.

That's a sure way to get caught. If you're working at burger king and you go out and buy a Bentley then people know that soimething's up and the government is gonna want to know how you bought it. It's smarter to put it away and spend slowly.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

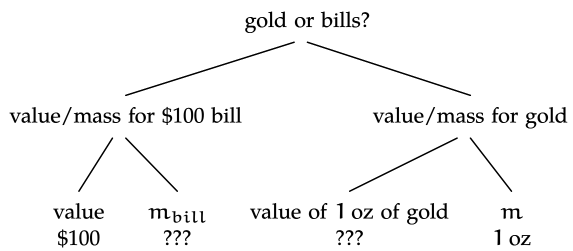

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

wouldnt it really just be weight? either quantities get very heavy at reasonable carrying volumes

Possibly, but I feel like it's important to take it into consideration for the sake of determining probability with more evidence. It wouldn't cost us much more estimation time, either.

Nevermind, I just realized that we completely disregard it in the next sentence.

This reminds me of a certain knapsack problem from 6.00.

Hehe yup

I think shape matters too. The gold may be more (or less) awkward to carry, thus allowing less (or more) weight to be carried.

But in the knapsack problem you have many possible items to choose from. Here it seems like we're choosing between all gold or all bills - a binary choice. And we seem to have unlimited supplies of both.

do people really care about the weight or volume? they can just easily save the bills in the bank right?

"Carrying capacity." As in, how much can you carry out of the bank when you're robbing it? Weight and volume are the important factors there.

When you say you are going to forget volume, do you mean you are simply going to assume that the backpack is big enough for it? This seems wrong because I feel that there would be a big difference between 50 lbs of gold and dollar bills.

Maybe it's because you are more likely to be limited by weight before you would be limited by volume.

Yes, I feel like volume is the limiting factor... bills are lighter and easier to transport but not if you have to carry 20 backpacks' worth.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

Having broken into a bank vault, should we take the bills or the gold?

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

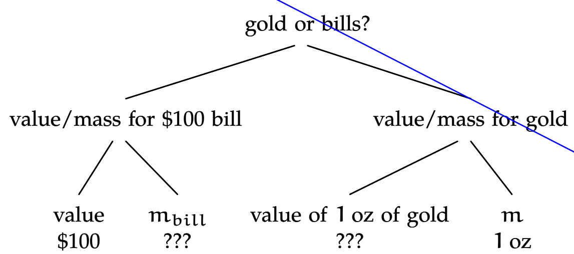

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

isn't this a rather large assumption? i understood this for the 747 CDs problem, but not as much for this one.

I think the idea here is that you would max out the mass you can possibly carry before maxing out the volume you can take.

The weight probably isn't an issue for the cash - but for the bars of gold I think it would be the limiting factor for sure.

It would be nice to go back after finding the best choice according to mass and check that for a substantial haul it doesn't violate a reasonable volume (like a truck).

I think it would be weight for the gold and volume for the money. How would we be able to separate them this way?

the comments on this reading are hilarious – of course MIT students would approach a bank robbery this way

How does the final answer change if you choose volume instead?

That's interesting...so presumably, you are limited in volume or mass. Since the density of gold >> density of paper, then here it seems like the mass is the limiting factor, and not volume. Also, gold is typically valued per pound, not per m^3. I kind of feel like per volume is just awkward.

I would assume that the answer to this would be pretty obvious. Gold density >> paper/wood density, but estimating the order of magnitude of difference might be an interesting problem.

Is this assuming we'll be able to take the same amount (mass) of either the bills or the gold? i.e we'll be able to take a knapsack full of either or something?

no the "knapsack" would assume a volume limit. instead think of it as having infinitely many bags and filling them until you can't carry any more.

Yeah this is pretty clever. i mean it makes sense but i wouldn't have thought of this–if i was robbing a bank for some reason i'd probably just grab whatever i could take. but here he's saying that if you can only carry a certain amount (weight), which one would end up giving you more value for the same mass.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

Having broken into a bank vault, should we take the bills or the gold?

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

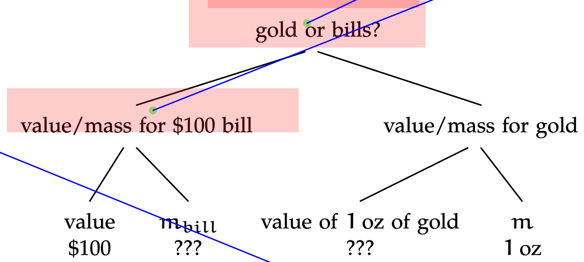

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

My intuition would say that volume is the limiting factor.. if we are carrying the cash/gold out in a duffle bag, it seems that we could in theory fit as much weight as we want but the volume of the duffle is what is going to limit exactly how much we can take.

I like that this chart is here- it's very clear and the labels help to show exactly what we know and what we're looking for.

I think that you should specify before this diagram that we are talking about 100$ bills

Or maybe mention that we'll be determining what order to take bills versus gold (i.e., what order of $100 bills, $50 bills, $20 bills, gold, etc.), so we'll measure the value/mass of a N dollar bill with value $N since all bills have the same mass.

Solving for an arbitrary bill value does seem like a good idea for a more detailed analysis.

Yeah, might banks have greater denominations anyway?

Never mind - apparently they stopped printing anything above $100 in 1969!

Whoops – the previous version of this section did talk about $100 bills. But I just changed it to be more general (so that it better fits with the use of plausible ranges). But I forgot to remove from the tree the reference to $100 bills.

In 6.00 last semester we had to write an algorithm to solve a similar problem. It was more simplified, involving small point values instead of actual money, so I'm very interested in seeing how this works out.

This question is actually much different. In 6.00, we dealt with various items of weight and value and had to use power sets to decide with set was the best. In this situation, we are assuming unlimited amounts of gold and cash...the question being which one is the better value per weight or mass.

This is such a simple and nice example for the tree. You could introduce something like this when you first introduce divide and conquer.

It seems like it's been a while since we had a tree diagram, and coming back to them makes me realize that I really really like them as a way to organize information.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

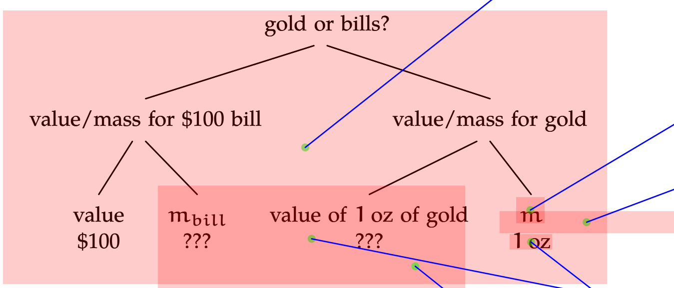

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

At first glance, I couldn't imagine gold even coming close in value per mass. Now that I'm thinking about it though, small amounts of gold are worth a lot of money. This is really interesting.

This is an excellent d&a diagram for this problem.

agreed

Divide and conquer at its best!

agreed...if you throw in volume there you might even be able to figure out which is easier to get out of the bank!

probably want to label mgold since other one is labeled mbill.

i just realized why it only says m. regardless, it would nice to have consistency from both sides.

by m you mean weight, not mass correct?

I think it means mass - oz is a measure of mass.

an oz is actually a measure of weight, since 16 ounces make a pound, and a pound is a measurement of weight, technically.

But he later converts ounces into grams, which is a measure of mass. So we do in fact want to find the mass of the gold, so it's correct. He uses ounces here out of convenience because gold is priced in ounces.

I like that this tree immediately demonstrates that our unknown quantities are opposites (not literally) in each case, and that therefore we will have to do two distinct estimations. (Yeah, I know it's trivial, but I still liked it.)

Me too, it's funny how we know very well the opposite parameters for gold and paper money. I guess it speaks to the disconnection of paper currency from real material value.

yeah I agree, this tree is really helpful with the question marks to highlight what the unknown values are.

wouldn't grams work better for estimation? the whole powers of 10 thing....?

Perhaps, but the value of gold is most often referred to in terms of $/ounce.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

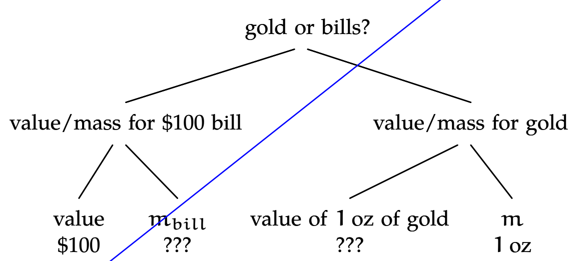

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

perhaps branches?

What are these again?

I think I means making a best guess.

I think these are single-number estimates, instead of giving it a range.

that sounds right although i dont remember seeing this terminology before.

I agree, but didn't we use the plausible range method with a range? Maybe I am jumping to the second round of analysis too soon.

Yeah, that's what I would have chosen to do first too.

So do we set the probability of the plausible range to 2/3 again here?

seems low to me

well it's more like 1200/oz but i guess it's in the same order of magnitude.

are the fluctuations in the value of gold greater than that of bills?

well I would think that the fluctuations would be similar. If the value of gold went up then the dollar went down. In terms of goods gold is probably more stable

The fluctuation of gold is supposed to be zero. It's the bills that are supposed to fluctuate.

what if we don't remember what it is at all?

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

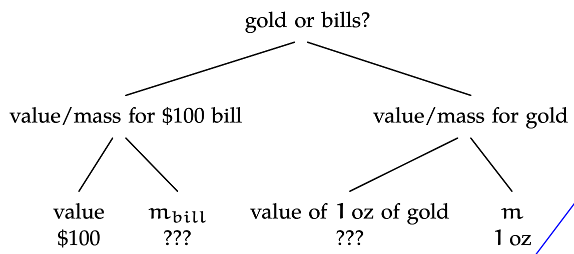

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

i feel like a lot of the problems in this book are done based on facts you "vaguely remember" but that i have no familiarity with whatsoever. the closest i could come to questimating this would be knowing how much some pieces of jewelry cost, but even then, i have no idea how much of it is for the gold and how much is for labor/advertising, etc.?

The price of gold is often mentioned in news and business reports on the radio and in newspapers, so it is something that people can at least know to an order of magnitude. Your idea about jewelry is a good one for bounding the price/oz. if you had to estimate.

Except that I know that carat matters when discussing gold, and I have no idea what that means for either the gold bars or the price of gold jewelery, as I've never seen the price for 24-carat samples of either.

So, when you do the calculation for this problem, how do you incorporate the price elasticity of gold?

I personally have no background in this area, I probably would have made this estimating using what I know about the price of gold jewelry (estimating the cost of a ring, the mass of that ring, and then making the proper conversions). With this method I actually got $1000 which is pretty close!

I saw that you know a lot of little stories (historical evidence or every day life evidence) that help you tremendously in your estimation

these stories are probably also key in helping to actually remembering the actual values of the estimates

Even if I knew the historical evidence, i still wouldnt know if it was correct or not to guess $800/oz. Would it have been a problem if I would have guessed anything from $300 - $1200?

I think that if you took the average value of your range and used it per oz, you would be fine.

Yes, it is quite entertaining when he launches into detailed anecdotes but not quite applicable to my problem solving process...

I think that we will accumulate more stories like these throughout our lifetimes too.

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

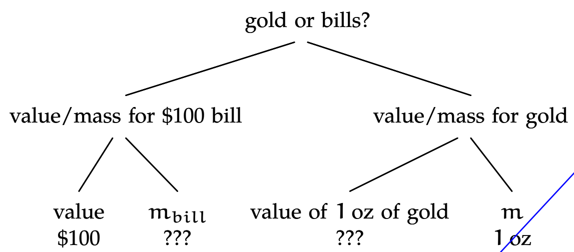

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

I enjoy reading annodotes like this because I am able to relate things unrelated to mathematics, physics to approximation. It makes everything seem more real and useful.

I agree. It makes the process feel more intuitive.

kind of a random piece of knoweldge to know

excellent anecdote, but is it expected we will develop such a large repertoire of anecdotes to help ourselves out?

I just saw this yesterday on The Pacific. I was wondering why it seemed so cheap

I'm amazed that you know this bit of history! And yet, not the price of a current-day gold piece

Right, if I had to look this up - might as well look up price of gold currently. I would try to estimate how much my last purchase of gold weighed and how much it cost instead.

Where did you get $35? Or is that a historical fact that you used because you knew it. If we weren't so gifted historically, how would we get an estimate?

Oh nevermind. I guess no matter what the estimate, if we could figure it out to be around $50 or so, it puts you in the ballpark of your $800 estimate.

But I still do feel that we have to rely a lot on previous knowledge or fact.

I'm way less likely to know anything about this than the current value of gold, which I also do not know. I'd be in trouble on this problem.

I don't think he is suggesting this as the "right" or only way of going about estimating this. Since he seems to have a great memory of history then this works for him. I personally think trying to remember the weight and price of a recent gold purchase is another fine approach.

As someone that has never researched the price of gold or purchased gold in the recent past I wouldn't have any clue where to start. I guess after doing the reading I could start with a guess of $800/oz though

8.2.2 Gold or bills?

The next estimation example is dedicated to readers who forgot careers in the financial industry for less lucrative careers in teaching and research.

▶ *Having broken into a bank vault, should we take the bills or the gold?*

The answer depends partly on the ease and costs of fencing the loot – an analysis beyond the scope of this book. Within our scope is the following question: Which choice lets us carry out the most money? Our carrying capacity is limited by weight and volume. For this analysis, let's assume that the more stringent limit comes from weight or mass. Then the decision divides into two subproblems: the value per mass for US bills and the value per mass for gold. In order to decide which to take, we'll compute both values per mass and their respective plausible ranges.

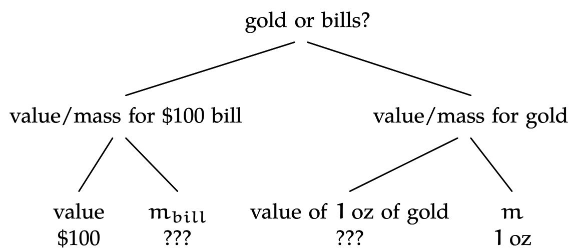

Two leaves have defined values: the value of a bill and the mass of 1 oz (1 ounce) of gold. The two other leaves need divide-and-conquer estimates. In the first round of analysis, make point estimates; then, in the second round, account for the uncertainty by using the plausible-range method of Section 8.2.

The value of gold is, I vaguely remember, around $800/oz. As a rough check on the value – for example, should it be $80/oz or $8000/oz? – here is a historical method. In 1945, at the end of World War 2, the British empire had exhausted its resources while the United States became the world's leading economic power. The gold standard, which fell apart during the depression, was accordingly reinstituted in terms of the dollar: $35 would be the value of 1 oz of gold. Since then, inflation has probably devalued the dollar by a factor of 10 or more, so gold should be worth around $350/oz. My vague memory of $800/oz therefore seems reasonable.

Hence all the commercials saying "trade your cash for gold!"

Not sure that this value is accurate anymore...I thought gold had been trading at more than $1000/oz for some time? $1000/oz is a nicer number to work with anyway.

Well like anything the price of gold goes up with demand, so all the commercials saying cash for gold are trying to cash in on the huge swing towards gold.

Yup, it's been trading at around $1000/oz. Since I haven't read till the end yet, I'm assuming there's a reason he's using 80/800/8000?

I agree - if I were doing this problem I would have chosen 100 - more accurate and easier to deal with

You mean $1000, not $100 right? Gold is almost at $1200 to be exact.

Yep, I looked it up too, just to be sure. It's pretty close to 1200$/oz.

strangely enough, if you say there is 4% inflation per year 35*(1.04)^80 = $806.74

so it's totally reasonable, but isn't 350 closer to 80 than 800?

I feel like I wouldn't know this kind of information off hand. I guess if I really wanted to make an accurate guess, I would google the value of gold before hand.

For the bill, its mass breaks into density (ρ) times volume (V), and volume breaks into width (w) times height (h) times thickness (t). To estimate the height and width, I could lay down a ruler or just find any bill – all US bills are the same size – and eyeball its dimensions. A $1 bill seems to be few inches high and 6 in wide. In metric units, those dimensions are $h \sim 6$ cm and $w \sim 15$ cm. [To improve your judgment for sizes, first make guesses; then, if you feel unsure, check the guess using a ruler to check. With practice, your need for the ruler will decrease and your confidence and accuracy will increase.]

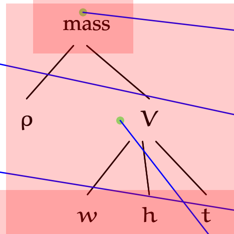

The thickness, alas, is not easy to estimate by eye or with a ruler. Is the thickness 1 mm or 0.1 mm or 0.01 mm? Having experience with such small lengths, my eye does not not help much. My ruler is calibrated in steps of 1 mm, from which I see that a piece of paper is significantly smaller than 1 mm, but I cannot easily see how much smaller.

An accurate divide-and-conquer estimate, we learned in Section 8.2, depends on replacing a vaguely understood quantity with accurately known quantities. Therefore to estimate the thickness accurately, I connect it to familiar quantities. Bills are made from paper, a ubiquitous substance (despite hype about the paperless office). Indeed, a ream of printer paper is just around the corner. The thickness of the ream and the number of sheets that it contains determines the thickness of one sheet:

$$t = \frac{t_{\text{ream}}}{N_{\text{ream}}}.$$

You might call this approach 'multiply and conquer'. The general lesson for accurate estimation is to magnify values much smaller than our experience, and and to shrink values much bigger than our experience.

The magnification argument adds one level to the tree and replaces one leaf with two leaves on the new level. Two of the five leaf nodes are already estimated. A ream contains 500 sheets ($N_{\text{ream}} = 500$) and has a thickness of roughly 2 in or 5 cm.

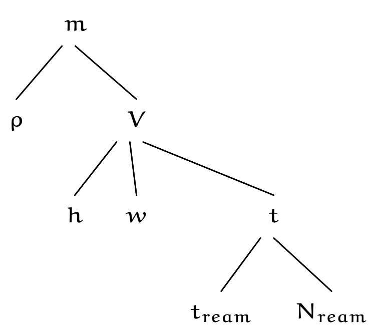

What is the estimate for ρ, the density of a bill?

Even though the information in this tree is obvious, I still like that it's here because it adds clarity.

If we are going to use the volume to calculate the mass, why don't we also consider the volume as part of the bigger problem.

what is a good way to estimate the thickness? since its so small

I don't think you would estimate the thickness of an individual bill. But what if you estimated the thickness of a stack of bills? You know that a stack of 500 pages of paper is 2-3 inches (thanks to Athena paper!). I would say the thickness of a bill is similar to the thickness of a sheet of paper. Does someone have a better idea?

I like how you created an additional tree instead of trying to cram this stuff onto the previous tree. It makes everything look more simple and clear.

If we could use a ruler on the bills, couldn't we have done the same for the piece of paper...

I don't think he's saying out should measure it, just to get an idea in your mind of how big it might be. The paper would be bigger then a standard ruler.

handy trick... a dollar bill is about 6" long. i've been known to use it as a ruler before.

This would be a more interesting question with a different currency, with different sized bills, like the Euro

Is it necessary to eyeball when you could use your body to measure. For instance, your thumb is probably about 1 inch from the tip to the closest knuckle. Why not measure the bill using that.

This reminds me of the reading yesterday, however, if we did all of that approximation, it would get pretty complicated.

For the bill, its mass breaks into density (ρ) times volume (V), and volume breaks into width (w) times height (h) times thickness (t). To estimate the height and width, I could lay down a ruler or just find any bill – all US bills are the same size – and eyeball its dimensions. A $1 bill seems to be few inches high and 6 in wide. In metric units, those dimensions are $h \sim 6\,\text{cm}$ and $w \sim 15\,\text{cm}$. [To improve your judgment for sizes, first make guesses; then, if you feel unsure, check the guess using a ruler to check. With practice, your need for the ruler will decrease and your confidence and accuracy will increase.]

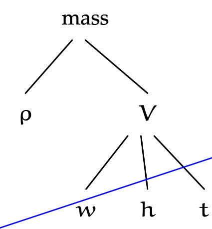

The thickness, alas, is not easy to estimate by eye or with a ruler. Is the thickness 1 mm or 0.1 mm or 0.01 mm? Having experience with such small lengths, my eye does not not help much. My ruler is calibrated in steps of 1 mm, from which I see that a piece of paper is significantly smaller than 1 mm, but I cannot easily see how much smaller.

An accurate divide-and-conquer estimate, we learned in Section 8.2, depends on replacing a vaguely understood quantity with accurately known quantities. Therefore to estimate the thickness accurately, I connect it to familiar quantities. Bills are made from paper, a ubiquitous substance (despite hype about the paperless office). Indeed, a ream of printer paper is just around the corner. The thickness of the ream and the number of sheets that it contains determines the thickness of one sheet:

$$t = \frac{t_{\text{ream}}}{N_{\text{ream}}}.$$

You might call this approach 'multiply and conquer'. The general lesson for accurate estimation is to magnify values much smaller than our experience, and and to shrink values much bigger than our experience.

The magnification argument adds one level to the tree and replaces one leaf with two leaves on the new level. Two of the five leaf nodes are already estimated. A ream contains 500 sheets ($N_{\text{ream}} = 500$) and has a thickness of roughly 2 in or 5 cm.

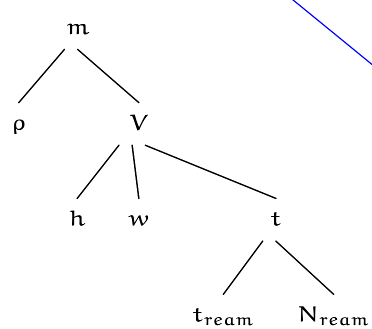

▶ *What is the estimate for ρ, the density of a bill?*

My favorite technique is to use various 'segments' of my fingers for measuring. I know the last segment of my index finger is an inch, the second to last 1.25", and various others. I also know that my hand spread is 8.5", for larger measurements.

This comes in handy since I think most people (myself included) tend to misjudge how big or small an inch actually is.

yeah, I think this is a great tool....just so long as you have stopped growing. Don't have a 13 year old try to memorize his dimensions. Although it might be easier to find a segment that is close to standard (1in, 2in, 10 cm, etc.) since 1.25 although its exact, might not be the easiest length to work with .

I was thinking it'd be easier to estimate a stack of bills, like you see in the money suitcases of movies! That's closer to what you'd pick up anyway.

who has experience with a stack of bills in a suitcase??

Maybe just guessing from seeing them in movies... I think the height and width are easy to estimate, but the depth is harder and it would probably be easier to do in stacks. Its hard to estimate when you get as small as mm.

Yeah i agree. Also, we could assume cash weighs about the same as printer paper. We know roughly how much a ream of paper weighs then we can divide that by about 6 to get a stack of cash.

This is an awesome estimation idea, I would have never thought of it myself-thanks!

good life skill!

Should we make guesses within a range first? I has been very helpful to me to make an estimate since some ranges just seem foolish

I feel liek this should have been said earlier in the book .. reiterating it is good, but i feel like this particular verbage would have been more useful earlier.

redundant

I like this suggestion.

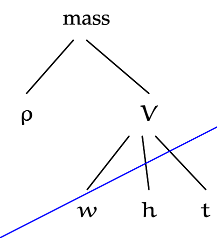

For the bill, its mass breaks into density (ρ) times volume (V), and volume breaks into width (w) times height (h) times thickness (t). To estimate the height and width, I could lay down a ruler or just find any bill – all US bills are the same size – and eyeball its dimensions. A $1 bill seems to be few inches high and 6 in wide. In metric units, those dimensions are $h \sim 6$ cm and $w \sim 15$ cm. [To improve your judgment for sizes, first make guesses; then, if you feel unsure, check the guess using a ruler to check. With practice, your need for the ruler will decrease and your confidence and accuracy will increase.]

The thickness, alas, is not easy to estimate by eye or with a ruler. Is the thickness 1 mm or 0.1 mm or 0.01 mm? Having experience with such small lengths, my eye does not not help much. My ruler is calibrated in steps of 1 mm, from which I see that a piece of paper is significantly smaller than 1 mm, but I cannot easily see how much smaller.

An accurate divide-and-conquer estimate, we learned in Section 8.2, depends on replacing a vaguely understood quantity with accurately known quantities. Therefore to estimate the thickness accurately, I connect it to familiar quantities. Bills are made from paper, a ubiquitous substance (despite hype about the paperless office). Indeed, a ream of printer paper is just around the corner. The thickness of the ream and the number of sheets that it contains determines the thickness of one sheet:

$$t = \frac{t_{\text{ream}}}{N_{\text{ream}}}.$$

You might call this approach 'multiply and conquer'. The general lesson for accurate estimation is to magnify values much smaller than our experience, and and to shrink values much bigger than our experience.

The magnification argument adds one level to the tree and replaces one leaf with two leaves on the new level. Two of the five leaf nodes are already estimated. A ream contains 500 sheets ($N_{\text{ream}} = 500$) and has a thickness of roughly 2 in or 5 cm.

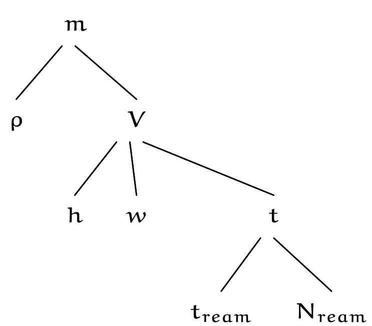

What is the estimate for ρ, the density of a bill?

This actually is true. I can usually accurately estimate length of up to 3 meters with 4-5 cm accuracy. Not too bad...

This reminds me about a lecture a couple weeks ago talking about "perfect practice" (the lecture about chess masters).

Would it be easier to measure the thickness of a stack of bills, and then divide by the number of bills in the stack? It seems like that would be a lot easier, unless there's some issue with that method that I'm not seeing.

I would think that unless you had a stack of freshly printed bills this method wouldn't work because the wrinkles and creases in used bills are probably much larger than the thickness of a bill. So a stack of bills would probably be 3 or 4 times the thickness it should be.

For the Mech E students, recall that a sheet of paper is about .003". You use this to touch off the tool tip w/o breaking it.

let me just whip out my calipers... (yeah, course 2!)

SHould this have despite?

Should this say "having no experience with such small lengths, my eye does not help much"?

Agreed.

you have an extra "not" in here

minor quibble– i have no idea what section 8.2 is. can you start naming the files with their section number rather than reading number so we can go back and reread them?

If he renamed them I think the text would get really wordy- since this is meant for people who will own this book, it will be really easy to flip back to section 8.2.

I really appreciate the reference and reiteration of this concept.

I agree. The succinctness was also great such that if you remember what was discussed in 8.2, you're not bombarded with a long review.

For the bill, its mass breaks into density (ρ) times volume (V), and volume breaks into width (w) times height (h) times thickness (t). To estimate the height and width, I could lay down a ruler or just find any bill – all US bills are the same size – and eyeball its dimensions. A $1 bill seems to be few inches high and 6 in wide. In metric units, those dimensions are $h \sim 6$ cm and $w \sim 15$ cm. [To improve your judgment for sizes, first make guesses; then, if you feel unsure, check the guess using a ruler to check. With practice, your need for the ruler will decrease and your confidence and accuracy will increase.]

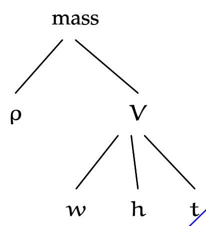

The thickness, alas, is not easy to estimate by eye or with a ruler. Is the thickness 1 mm or 0.1 mm or 0.01 mm? Having experience with such small lengths, my eye does not not help much. My ruler is calibrated in steps of 1 mm, from which I see that a piece of paper is significantly smaller than 1 mm, but I cannot easily see how much smaller.

An accurate divide-and-conquer estimate, we learned in Section 8.2, depends on replacing a vaguely understood quantity with accurately known quantities. Therefore to estimate the thickness accurately, I connect it to familiar quantities. Bills are made from paper, a ubiquitous substance (despite hype about the paperless office). Indeed, a ream of printer paper is just around the corner. The thickness of the ream and the number of sheets that it contains determines the thickness of one sheet:

$$t = \frac{t_{\text{ream}}}{N_{\text{ream}}}.$$

You might call this approach 'multiply and conquer'. The general lesson for accurate estimation is to magnify values much smaller than our experience, and and to shrink values much bigger than our experience.

The magnification argument adds one level to the tree and replaces one leaf with two leaves on the new level. Two of the five leaf nodes are already estimated. A ream contains 500 sheets ($N_{\text{ream}} = 500$) and has a thickness of roughly 2 in or 5 cm.

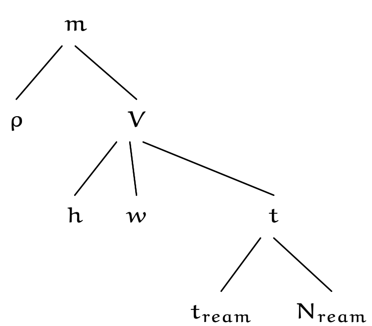

What is the estimate for ρ, the density of a bill?

hmm, so I thought they were made from cloth. I looked it up and it turns out they are part paper, and part cloth (cotton, silk and linen).

Bills really have silk in them? wow. Explains why the CAN go through the washer...

Lol, maybe to be more exact, we can say 2 bills is as thick as 3 sheets of paper. I'm not sure if this is necessary in an approximation class though.

wouild it not be easier to conform a number of bills to 1mm and count the number and deduce the thickness that way?

Is this the same ream that our paper comes in?

Yeah–obviously the bill paper isn't the exact same as our athena cluster printer paper, but for the sake of estimation it will give us a pretty close estimate.

Yeah, a ream is actually an official unit of measurement for paper - 500 sheets.

I wonder what the range of thicknesses of paper is- I bet their actually is a lot of variety

I was just thinking of this as a plausible way to figure it out

I enjoy the way you estimate using daily items, however, in cases like this, when we would have to go out to search for a ream of paper, it would seem that a better example might be possible.

As others have suggested I thought this was a bit involved and would've been easier if you stacked 1's to make 1mm and then figured out the thickness from that.

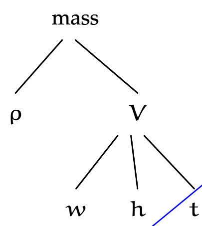

For the bill, its mass breaks into density (ρ) times volume (V), and volume breaks into width (w) times height (h) times thickness (t). To estimate the height and width, I could lay down a ruler or just find any bill – all US bills are the same size – and eyeball its dimensions. A \$1 bill seems to be few inches high and 6 in wide. In metric units, those dimensions are $h \sim 6\,\text{cm}$ and $w \sim 15\,\text{cm}$. [To improve your judgment for sizes, first make guesses; then, if you feel unsure, check the guess using a ruler to check. With practice, your need for the ruler will decrease and your confidence and accuracy will increase.]

The thickness, alas, is not easy to estimate by eye or with a ruler. Is the thickness 1 mm or 0.1 mm or 0.01 mm? Having experience with such small lengths, my eye does not not help much. My ruler is calibrated in steps of 1 mm, from which I see that a piece of paper is significantly smaller than 1 mm, but I cannot easily see how much smaller.

An accurate divide-and-conquer estimate, we learned in Section 8.2, depends on replacing a vaguely understood quantity with accurately known quantities. Therefore to estimate the thickness accurately, I connect it to familiar quantities. Bills are made from paper, a ubiquitous substance (despite hype about the paperless office). Indeed, a ream of printer paper is just around the corner. The thickness of the ream and the number of sheets that it contains determines the thickness of one sheet:

$$t = \frac{t_{\text{ream}}}{N_{\text{ream}}}.$$

You might call this approach 'multiply and conquer'. The general lesson for accurate estimation is to magnify values much smaller than our experience, and and to shrink values much bigger than our experience.

The magnification argument adds one level to the tree and replaces one leaf with two leaves on the new level. Two of the five leaf nodes are already estimated. A ream contains 500 sheets ($N_{\text{ream}} = 500$) and has a thickness of roughly 2 in or 5 cm.

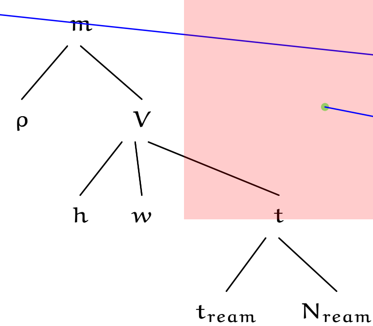

▶ *What is the estimate for ρ, the density of a bill?*

Above, I assumed we would figure out the thickness by stacking enough bills to measure a millimeter or two. I'm happy to see a version of my solution was eventually used.

Yes, me too. I just made a comment above asking if it would be easier to measure a stack of bills instead of just one. Looks like I should have just read a few paragraphs ahead!

Ha, I didn't think of this at all... I wanted to take it from the other approach of thinking of something comparably small (i.e., the width of a hair), but this makes more sense.

This is a good approach; one I didn't think about until I got here. For the above comment on stacking bills; bills get wrinkled and unless you ironed 100 bills and stacked them, compressing slightly to take out air, you won't get an accurate result. The paper in a ream is compressed and all nice for your ease of estimation.

It seems like most used bills take up significantly more space vertically when stacked due to wrinkles that develop through use. I don't think this approach would take this into account.

I feel like we have already learned this from d&a questions in the past. When we estimate things such as the fuel efficiency of a plane, we "multiply" by relating to a long plane ride in which we know the price of a ticket, # people on board and time of flight.

Clever. A really simple idea, but I would have never thought of putting a ream of bills together, measure their thickness, and divide by the number of bills.

I feel like this paragraph would also go better in the D&C section ... it's really well written

This is such an incredibly useful thing to know.

It's nice to point it out explicitly too, even if we have done in before.

Yea, I'd never really thought about it like that but its what we try to do all the time.

typo: repeated "and"

Since this is the same chart as at the top of this page, maybe for space efficiency you only need this one?

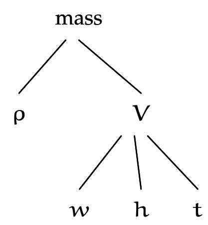

For the bill, its mass breaks into density (ρ) times volume (V), and volume breaks into width (w) times height (h) times thickness (t). To estimate the height and width, I could lay down a ruler or just find any bill – all US bills are the same size – and eyeball its dimensions. A $1 bill seems to be few inches high and 6 in wide. In metric units, those dimensions are $h \sim 6\,\mathrm{cm}$ and $w \sim 15\,\mathrm{cm}$. [To improve your judgment for sizes, first make guesses; then, if you feel unsure, check the guess using a ruler to check. With practice, your need for the ruler will decrease and your confidence and accuracy will increase.]

The thickness, alas, is not easy to estimate by eye or with a ruler. Is the thickness 1 mm or 0.1 mm or 0.01 mm? Having experience with such small lengths, my eye does not not help much. My ruler is calibrated in steps of 1 mm, from which I see that a piece of paper is significantly smaller than 1 mm, but I cannot easily see how much smaller.

An accurate divide-and-conquer estimate, we learned in Section 8.2, depends on replacing a vaguely understood quantity with accurately known quantities. Therefore to estimate the thickness accurately, I connect it to familiar quantities. Bills are made from paper, a ubiquitous substance (despite hype about the paperless office). Indeed, a ream of printer paper is just around the corner. The thickness of the ream and the number of sheets that it contains determines the thickness of one sheet:

$$t = \frac{t_{\mathrm{ream}}}{N_{\mathrm{ream}}}.$$

You might call this approach 'multiply and conquer'. The general lesson for accurate estimation is to magnify values much smaller than our experience, and and to shrink values much bigger than our experience.

The magnification argument adds one level to the tree and replaces one leaf with two leaves on the new level. Two of the five leaf nodes are already estimated. A ream contains 500 sheets ($N_{\mathrm{ream}} = 500$) and has a thickness of roughly 2 in or 5 cm.

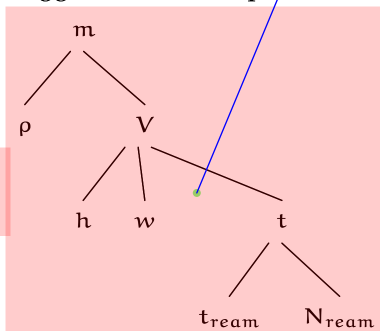

What is the estimate for ρ, the density of a bill?

You might want to do the division out for the thickness of an individual sheet of paper here as it feels like you kind of just leave us hanging.

I don't think it is necessary as he has written the formula for it above.

So we just assumed that the thickness of printer paper is the same as cash?

this gets you an answer surprisingly close to the actual thickness of paper - as someone who does a fair amount of machining, I know regular printer paper is about .003" thick.

I find this chart more confusing than the others.

The only missing leaf value is ρ, the density of a bill. Connect this value to what you already know such as the densities of familiar substances. Bills are made of paper, whose density is hard to guess directly. However, paper is made of wood, whose density is easy to guess! Wood barely floats so its density is roughly that of water: $1\,\text{g}\,\text{cm}^{-3}$. Therefore the density of a bill is roughly $1\,\text{g}\,\text{cm}^{-3}$.

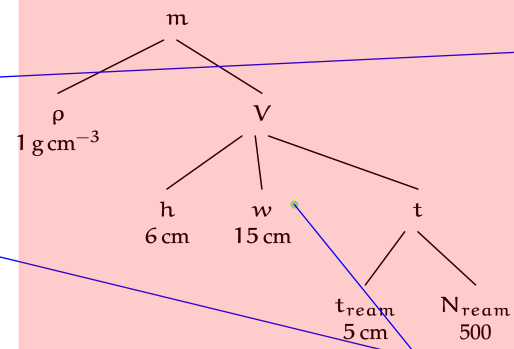

Now propagate the leaf values upward. The thickness of a bill is roughly 10^{-2} cm, so the volume of a bill is roughly

$$V \sim 6\,\text{cm} \times 15\,\text{cm} \times 10^{-2}\,\text{cm} \sim 1\,\text{cm}^3.$$

The mass is therefore

$$m \sim 1\,\text{cm}^3 \times 1\,\text{g}\,\text{cm}^{-3} \sim 1\,\text{g}.$$

and the value per mass of an $$N$ bill is therefore $$N/\text{g}$. How simple!

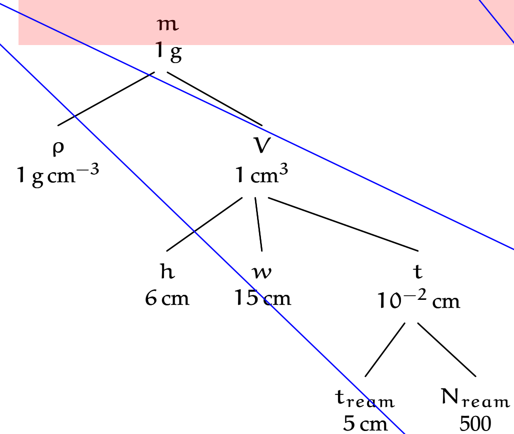

To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately the price of gold is usually quoted in dollars per ounce rather than dollars per gram, so my vague memory of $800/oz needs to be converted into metric units. One ounce is roughly 28 g; if the price of gold were $840/oz, the arithmetic is simple enough to do mentally, and produces $30/g. An exact division produces the slightly lower figure of $28/g. The result of this calculation is as follows: In the bank vault, first collect all the $100 bills that we can carry. If we have spare capacity, collect the $50 bills, the gold, and only then the $20 bills.

This order depends on the accuracy of the point estimates and would change if the estimates are significantly inaccurate. But how accurate are they? To analyze the accuracy, make plausible ranges for the leaf nodes

Obviously bills are different than paper, how much of a change would that create in the approx answer versus the actual?

I agree, I think the breakdown here is informative but when carrying out as many dollar bills as possible, a small variation in density could be crucial to the answer.

Are cotton and silk really that much denser than paper? I don't think they necessarily are, so it wouldn't be that bad to estimate them as wood. However, I don't know the density of gold off the top of my head and that might be harder!

Agreed – cotton and silk seem like the would also both barely float, giving them densities similar to that of wood.

Dollar bills are actually made from a mixture of cotton and linen. But this estimate should be pretty close still–all plant material.

and wood pulp/paper. http://www.ehow.com/how-does_4613179_what-dollar-bill-made.html

don't believe the first yahoo! answers that comes up when you search "what are bills made of"

Yep, I had the same thought; bills are not made of wood.

I like these figures. They make the writing very clear.

While it makes sense, I'm not that comfortable using this method. I don't really like it.

I agree - I think paper is pretty removed from wood in my mind...

If you don't want to trust this, weight a ream of paper and divide by its volume. I haven't done it, by my guess is that it checks out.

another way to think of it is that the weight of a ream of paper is about the same weight as a similar sized piece of wood. so a ream of paper is pretty heavy, but if you think about it, the same volume of pure water would probably be heavier than a ream of printer paper.

Aren't American bills made from cotton fiber paper? I'm not sure if this type of paper has the same density as wood.

The only missing leaf value is ρ, the density of a bill. Connect this value to what you already know such as the densities of familiar substances. Bills are made of paper, whose density is hard to guess directly. However, paper is made of wood, whose density is easy to guess! Wood barely floats so its density is roughly that of water: $1\,\mathrm{g\,cm^{-3}}$. Therefore the density of a bill is roughly $1\,\mathrm{g\,cm^{-3}}$.

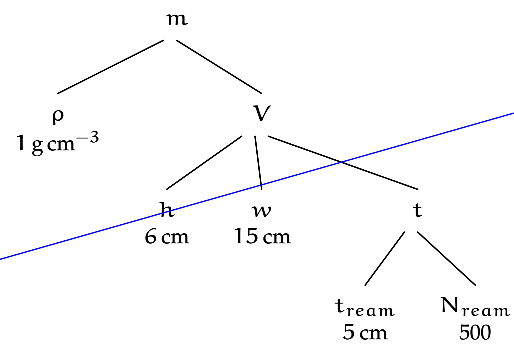

Now propagate the leaf values upward. The thickness of a bill is roughly 10^{-2} cm, so the volume of a bill is roughly

$$V \sim 6\,\mathrm{cm} \times 15\,\mathrm{cm} \times 10^{-2}\,\mathrm{cm} \sim 1\,\mathrm{cm}^3.$$

The mass is therefore

$$m \sim 1\,\mathrm{cm}^3 \times 1\,\mathrm{g\,cm^{-3}} \sim 1\,\mathrm{g}.$$

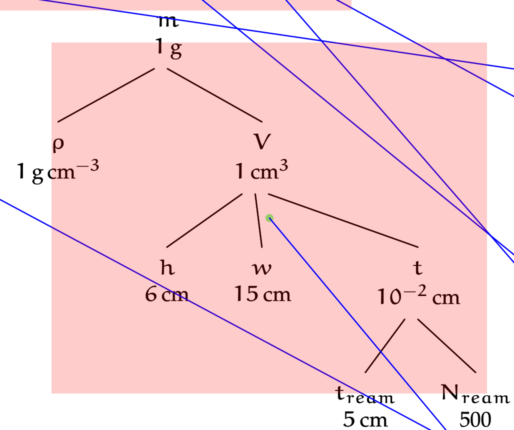

and the value per mass of an $\$N$ bill is therefore $\$N/\mathrm{g}$. How simple!

To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately the price of gold is usually quoted in dollars per ounce rather than dollars per gram, so my vague memory of \$800/oz needs to be converted into metric units. One ounce is roughly 28 g; if the price of gold were \$840/oz, the arithmetic is simple enough to do mentally, and produces \$30/g. An exact division produces the slightly lower figure of \$28/g. The result of this calculation is as follows: In the bank vault, first collect all the \$100 bills that we can carry. If we have spare capacity, collect the \$50 bills, the gold, and only then the \$20 bills.

This order depends on the accuracy of the point estimates and would change if the estimates are significantly inaccurate. But how accurate are they? To analyze the accuracy, make plausible ranges for the leaf nodes

What also floats in water? A duck. Exactly! So, logically... If... she.. weighs the same as a duck, she's made of wood. And therefore–? A witch

Monty Python would disagree with you here.!

Sadly, I had the same thought when I read that phrase. Also, not all wood floats, for those amateur boat-builders out there.

you could also get this guess by noticing that paper floats in water and thus is less than or equal to the density of water.

True. I guess we could reason that loggers use rivers to transport large pieces of wood, so the wood used in paper floats..

barely? I always thought wood floated pretty well. I mean, it can float with extra weight ontop of it even. (ie: a raft)

really interesting and cool way of approximating the density of a bill by relating paper to wood and considering the fact that it floats on water.

I understand this is a close approx. but why not use something like .9? It will likely play out in the long run as a significant mass when we multiply by thousands of bills.

It will still only be 10% different, and like we have been talking about in class, a factor of 1.1 is mostly going to disappear in combining uncertainties. Are you really sure it's closer to 0.9 than 0.98?

This should stick with the formatting of the rest of the paragraph, imo.

This is really cool. Though I feel that paper floats more easily than wood.

Another fun way to do this would be to fold a bill into a tiny box, and estimate its volume. Mine (it was a flat box) was 3mm*15mm*22mm=1.4cm^3. Pretty close, especially considering it's not perfectly compressed.

That's really clever! I wouldn't have thought to try that.

these figures really help me understand what is going on

The only missing leaf value is ρ, the density of a bill. Connect this value to what you already know such as the densities of familiar substances. Bills are made of paper, whose density is hard to guess directly. However, paper is made of wood, whose density is easy to guess! Wood barely floats so its density is roughly that of water: $1\,\mathrm{g\,cm^{-3}}$. Therefore the density of a bill is roughly $1\,\mathrm{g\,cm^{-3}}$.

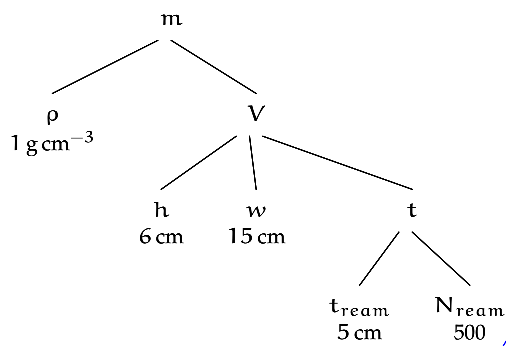

Now propagate the leaf values upward. The thickness of a bill is roughly 10^{-2} cm, so the volume of a bill is roughly

$$V \sim 6\,\mathrm{cm} \times 15\,\mathrm{cm} \times 10^{-2}\,\mathrm{cm} \sim 1\,\mathrm{cm}^3.$$

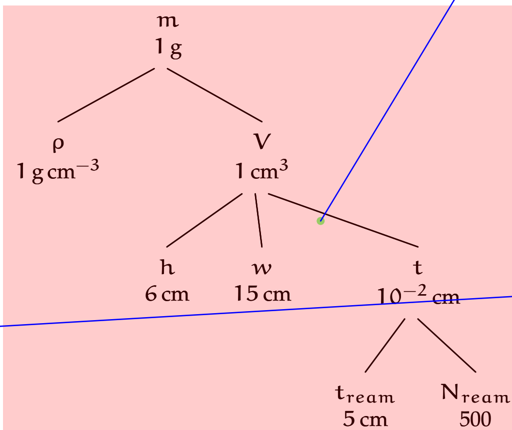

The mass is therefore

$$m \sim 1\,\mathrm{cm}^3 \times 1\,\mathrm{g\,cm^{-3}} \sim 1\,\mathrm{g}.$$

and the value per mass of an \$$N$ bill is therefore \$$N$/g. How simple!

To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately the price of gold is usually quoted in dollars per ounce rather than dollars per gram, so my vague memory of \$800/oz needs to be converted into metric units. One ounce is roughly 28 g; if the price of gold were \$840/oz, the arithmetic is simple enough to do mentally, and produces \$30/g. An exact division produces the slightly lower figure of \$28/g. The result of this calculation is as follows: In the bank vault, first collect all the \$100 bills that we can carry. If we have spare capacity, collect the \$50 bills, the gold, and only then the \$20 bills.

This order depends on the accuracy of the point estimates and would change if the estimates are significantly inaccurate. But how accurate are they? To analyze the accuracy, make plausible ranges for the leaf nodes

This might be a personal thing, but I feel these trees are a bit skewed and stretched out of place with respect to the text. Any way of reformatting them?

They do seem to be a bit funny since the branches are rather one sided, but I don't think it's problematic to the point of being detracting.

They could be more compact, but I like how there is an illustration for this problem considering we just dealt with a lot of known values.

I thought we were going to change to having sideways trees with the method of acting on them written on the branch? (multiplication, addition, all that stuff)

I plan on trying that. I still need to rewrite the tree-language-to-pdf compiler to generate trees that grow horizontally. Once that happens, all trees will grow sideways (and I'll probably need to check that they aren't too wide).

I agree. The diagram is useful in keeping track of all the values we've just estimated. Also it's nice to come full circle back to material from the beginning of the class, and the diagram helps to reinforce that idea.

i just tried to guess the mass of of a bill and came to this number. i feel like we can guess this, from things like knowing the mass of paperclips, pretty well. at least personally, i'm better equipped to guess the mass of a dollar bill than it's density and thickness.

I probably would have guessed that too from the very beginning. But since this is such an easy number to remember, I'll never have to do the calculations!

I think I'd personally hesitate when guessing something with a small mass just because I'm not really attuned to what a gram feels like (even versus a paperclip, cuz it has a higher density).

My first intuition was that it would be smaller than a gram, so I would have been off by at least a factor of 2 here if we hadn't done out all the calculations.

Having held a ream of paper, it seems odd to me that it would float, even though I looked up the density of paper and found that it is quite close to this estimate.

Try thinking about picking up a log of wood. It also floats but is pretty heavy....

The only missing leaf value is ρ, the density of a bill. Connect this value to what you already know such as the densities of familiar substances. Bills are made of paper, whose density is hard to guess directly. However, paper is made of wood, whose density is easy to guess! Wood barely floats so its density is roughly that of water: $1\,\mathrm{g\,cm^{-3}}$. Therefore the density of a bill is roughly $1\,\mathrm{g\,cm^{-3}}$.

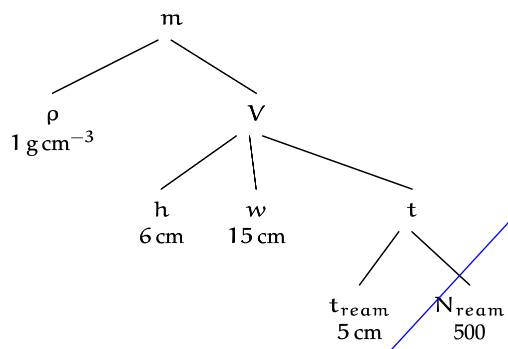

Now propagate the leaf values upward. The thickness of a bill is roughly 10^{-2} cm, so the volume of a bill is roughly

$$V \sim 6\,\mathrm{cm} \times 15\,\mathrm{cm} \times 10^{-2}\,\mathrm{cm}$$
$$\sim 1\,\mathrm{cm}^3.$$

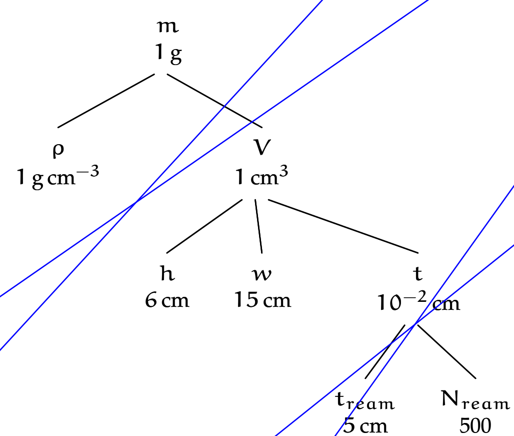

The mass is therefore

$$m \sim 1\,\mathrm{cm}^3 \times 1\,\mathrm{g\,cm^{-3}} \sim 1\,\mathrm{g}.$$

and the value per mass of an $N bill is therefore $N/g. How simple!

To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately the price of gold is usually quoted in dollars per ounce rather than dollars per gram, so my vague memory of $800/oz needs to be converted into metric units. One ounce is roughly 28 g; if the price of gold were $840/oz, the arithmetic is simple enough to do mentally, and produces $30/g. An exact division produces the slightly lower figure of $28/g. The result of this calculation is as follows: In the bank vault, first collect all the $100 bills that we can carry. If we have spare capacity, collect the $50 bills, the gold, and only then the $20 bills.

This order depends on the accuracy of the point estimates and would change if the estimates are significantly inaccurate. But how accurate are they? To analyze the accuracy, make plausible ranges for the leaf nodes

Grams and ounces in the same problem? i see comparison problems here.

I think the problem is that the price of gold is often specified in terms of USD/oz instead of USD/g (since the US uses imperial units). Converting shouldn't be too much of a problem, just another step or so.

Amazingly yes! This problem was a nice break from the drag concepts and a good refresher on dividing and conquering. Especially since so many of the unkowns were tricky, at least for a bit, to relate to a known fact.

vague/unclear what it refers to. Compare the value of per mass of dollars to the value per mass of gold.

I don't really see how this is vague or unclear... I thought it read fine.

This might be a slightly annoying, trivial comment, but I personally find arithmetic done in the paragraphs, even if it is really simple, harder to read than if it had just been split off.

Not something I can say is in my immediate memory.

I was trying to figure out how i would get this not knowing it... I think I would use measuring cups in my kitchen. I know they have both cups and liters on them, and I know how many ounces in a cup and how to convert from liters to grams for water.

I always find it easier to remember that 3.5 oz = 100g. And I know that from working with yarn...

Good to know. It does depend on the material tho right? Ounce is a volume while grams are mass?

Nice fudging to make the calculation easier.

You say that the value of gold is about $28/g, and the value of a $20 bill is $20/g, and yet you still say to take the $20 bills before the gold. Contradiction?

He says take the gold before the $20 bills...

The only missing leaf value is ρ, the density of a bill. Connect this value to what you already know such as the densities of familiar substances. Bills are made of paper, whose density is hard to guess directly. However, paper is made of wood, whose density is easy to guess! Wood barely floats so its density is roughly that of water: $1\,\text{g}\,\text{cm}^{-3}$. Therefore the density of a bill is roughly $1\,\text{g}\,\text{cm}^{-3}$.

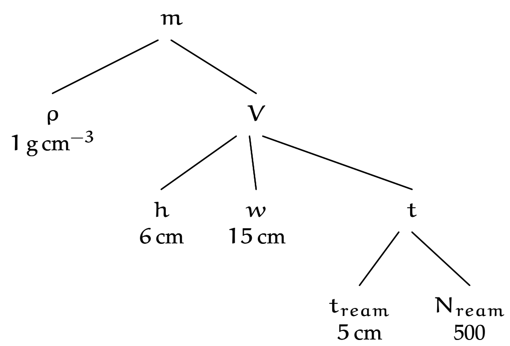

Now propagate the leaf values upward. The thickness of a bill is roughly 10^{-2} cm, so the volume of a bill is roughly

$$V \sim 6\,\text{cm} \times 15\,\text{cm} \times 10^{-2}\,\text{cm} \sim 1\,\text{cm}^3.$$

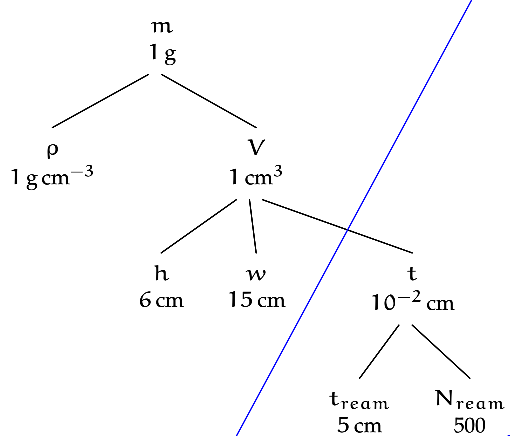

The mass is therefore

$$m \sim 1\,\text{cm}^3 \times 1\,\text{g}\,\text{cm}^{-3} \sim 1\,\text{g}.$$

and the value per mass of an $\$N$ bill is therefore $\$N/\text{g}$. How simple!

To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately the price of gold is usually quoted in dollars per ounce rather than dollars per gram, so my vague memory of \$800/oz needs to be converted into metric units. One ounce is roughly 28 g; if the price of gold were \$840/oz, the arithmetic is simple enough to do mentally, and produces \$30/g. An exact division produces the slightly lower figure of \$28/g. The result of this calculation is as follows: In the bank vault, first collect all the \$100 bills that we can carry. If we have spare capacity, collect the \$50 bills, the gold, and only then the \$20 bills. This order depends on the accuracy of the point estimates and would change if the estimates are significantly inaccurate. But how accurate are they? To analyze the accuracy, make plausible ranges for the leaf nodes

Good to keep in mind..

This problem is really counter-intuitive, most of us (including myself) would have initially said take the gold...quite an interesting and useful conclusion though!!

I wonder which is actually easier to launder. I wouldn't have the slightest idea how to launder gold bricks, but I feel like they could track the serials on 100s.

I was just thinking the same thing..the police can definitely track the serials fairly easily, whereas it think it's pretty easy to sell gold on the black market. You'd probably be in with that crowd if you were robbing a bank.

you could melt it down.

also...hahaha i had totally not thought about the different value amounts of the bill....hahaha

I think if you're stealing that many \$100 bills, it wouldn't be difficult for you to disperse it overseas, or to use it as "credit" without anyone actually spending it publicly.

Especially if you don't go and spend all of the bills on one big thing.

This is in interesting question with unexpected results. Thanks for sharing...although hopefully I will never need it.

So even if gold is valued at \$1000/oz as someone said before, then it doesn't change our answer because 1000/28 <50.

the current price of gold is 1100 per ounce.

So it still makes more sense to grab \$100 bills. Interesting! I wonder if this changes if we know someone who can turn the gold into jewelry, because per ounce i'm sure that is far more valuable, although there are time-hours in there too

Yeah, this is an interesting result. I wonder if the result would be the same if we chose to consider volume as the limiting factor, and not the mass. I'm guessing the higher bills would still win out since a gold brick is so large.

The only missing leaf value is ρ, the density of a bill. Connect this value to what you already know such as the densities of familiar substances. Bills are made of paper, whose density is hard to guess directly. However, paper is made of wood, whose density is easy to guess! Wood barely floats so its density is roughly that of water: $1\,\mathrm{g\,cm^{-3}}$. Therefore the density of a bill is roughly $1\,\mathrm{g\,cm^{-3}}$.

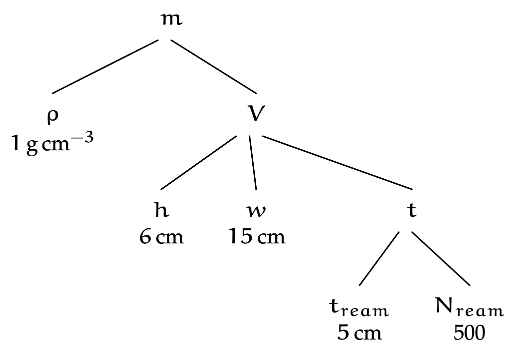

Now propagate the leaf values upward. The thickness of a bill is roughly 10^{-2} cm, so the volume of a bill is roughly

$$V \sim 6\,\mathrm{cm} \times 15\,\mathrm{cm} \times 10^{-2}\,\mathrm{cm} \sim 1\,\mathrm{cm}^3.$$

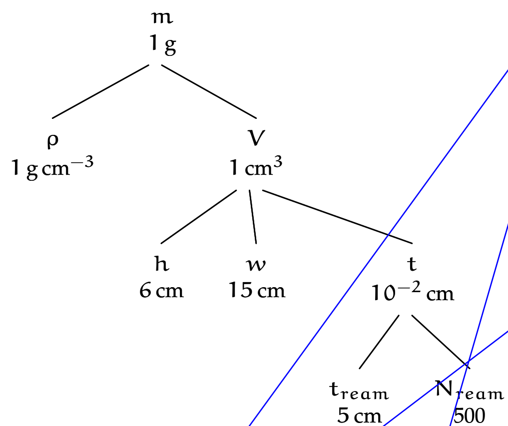

The mass is therefore

$$m \sim 1\,\mathrm{cm}^3 \times 1\,\mathrm{g\,cm^{-3}} \sim 1\,\mathrm{g}.$$

and the value per mass of an $\$N$ bill is therefore $\$N/\mathrm{g}$. How simple!

To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately the price of gold is usually quoted in dollars per ounce rather than dollars per gram, so my vague memory of \$800/oz needs to be converted into metric units. One ounce is roughly 28 g; if the price of gold were \$840/oz, the arithmetic is simple enough to do mentally, and produces \$30/g. An exact division produces the slightly lower figure of \$28/g. The result of this calculation is as follows: In the bank vault, first collect all the \$100 bills that we can carry. If we have spare capacity, collect the \$50 bills, the gold, and only then the \$20 bills.

This order depends on the accuracy of the point estimates and would change if the estimates are significantly inaccurate. But how accurate are they? To analyze the accuracy, make plausible ranges for the leaf nodes

Is this assuming we are only grabbing bills? Because otherwise wouldn't it make more sense to grab the 100's, 50's, and then some amount of gold?

also do banks store \$1000 dollar bills? or would these be way too easy to track?

That's what he's doing here–grabbing the 100's then the 50's then the gold then the 20's. the 20's have the least value per mass as shown by his calculations for bill values/mass

I'm surprised that you would even take the \$50 bills before the gold!

one thing to think about is that 20 dollar bills are far more widely circualted than 100 dollar bills. as a result depending on the weight that is determined that one can carry it may be more advisable to steal the 20 dollar bills first as they are harder to trace. just a thought..

That's adding in the problem of what to do with the stolen goods once we have them, which is outside the scope of this class and would probably get Sanjoy in trouble if he published it...

This seems really verbose. The word accurate/accuracy/inaccurate in used in these three sentences 4 times.

and propagate them upward – thereby obtaining plausible ranges for the value per mass of bills and gold.

Problem 8.2 Your plausible ranges
What are your plausible ranges for the five leaf quantities t_{ream}, N_{ream}, w, h, and ρ? Propagate them upward to get plausible ranges for the interior nodes including for the root node m.

Here are my ranges along with a few notes on how I estimated a few of them:

1. thickness of a ream, t_{ream}: $4\ldots6$ cm.
2. number of sheets in a ream, N_{ream}: 500. I'm almost certain that I remember this value correctly, but to be certain I confirmed it by looking at a label on a fresh ream.
3. width of a bill, w: $10\ldots20$ cm. A reasonable length estimate seemed to be 6 in but I could give or take a couple inches. In metric units, $4\ldots8$ in becomes (roughly) $10\ldots20$ cm.
4. height of a bill, h: $5\ldots7$ cm.
5. density of a bill, ρ: $0.8\ldots1.2\,\mathrm{g\,cm^{-3}}$. The argument for $\rho = 1\,\mathrm{g\,cm^{-3}}$ – that a bill is made from paper and paper is made from wood – seems reasonable. However, the many steps required to process wood into paper may reduce or increase the density slightly.

Now propagate these ranges upward. The plausible range for the thickness t becomes $0.8\ldots1.2\cdot10^{-2}$ cm. The plausible range for the volume V becomes $0.53\ldots1.27\,\mathrm{cm^3}$. The plausible range for the mass m becomes $0.50\ldots1.30$ g. The plausible range for the value per mass is $79...189/g (with a midpoint of $122/g).

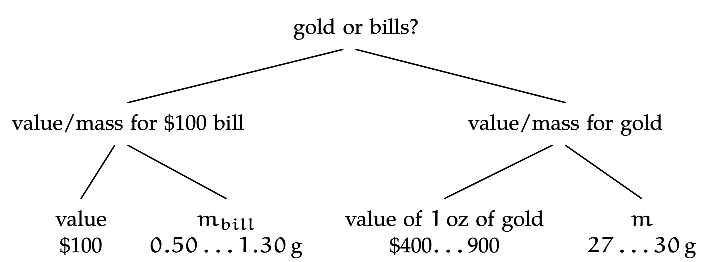

The next estimate is the value per mass of gold. I can be as accurate as I want in converting from ounces to grams. But I'll be lazy and try to remember the value while including uncertainty to reflect the

I really like that this is here!! It allows for your above conclusion to be so much more concise, but for the people who want to see how you came to it, there is this section to look over. Really great idea!! I think this set-up would do well in other sections too.

How do we determine how much weight to assign to volume as opposed to mass?

why do you not take into account that $$ paper is about twice as thick as standard printer paper??

Thats just as easy as looking up a value on the internet for something you don't know, so is there a rule like "if its on the label it doesn't count" for estimating stuff?

I was wondering about this too. I know he made the comment about paper reams being everywhere, but still it seems like cheating. How do you make the distinction between types of easy-of-access information?

Well, a ream is an official unit of measurement for paper - it doesn't change from company to company or label to label.

that is a super long bill

I'm a little confused here. It says the width of a bill is 10-20cm, as well as the length. This is a mistake right?

I feel like this is a particularly huge range for a bill that you (could) have right in front of you.

(To make my estimate I folded it in half and it appeared to be very close to 3 inches.)

Very smart, folding the bill in half would increase our accuracy.

would it be better to say length and width of a bill? since bills don't stand up?

and propagate them upward – thereby obtaining plausible ranges for the value per mass of bills and gold.

Problem 8.2 Your plausible ranges
What are your plausible ranges for the five leaf quantities t_{ream}, N_{ream}, w, h, and ρ? Propagate them upward to get plausible ranges for the interior nodes including for the root node m.

Here are my ranges along with a few notes on how I estimated a few of them:

1. thickness of a ream, t_{ream}: 4...6 cm.
2. number of sheets in a ream, N_{ream}: 500. I'm almost certain that I remember this value correctly, but to be certain I confirmed it by looking at a label on a fresh ream.
3. width of a bill, w: 10...20 cm. A reasonable length estimate seemed to be 6 in but I could give or take a couple inches. In metric units, 4...8 in becomes (roughly) 10...20 cm.
4. height of a bill, h: 5...7 cm.
5. density of a bill, ρ: 0.8...1.2 g cm^{-3}. The argument for $\rho = 1$ g cm^{-3} – that a bill is made from paper and paper is made from wood – seems reasonable. However, the many steps required to process wood into paper may reduce or increase the density slightly.

Now propagate these ranges upward. The plausible range for the thickness t becomes $0.8 \ldots 1.2 \cdot 10^{-2}$ cm. The plausible range for the volume V becomes 0.53...1.27 cm^3. The plausible range for the mass m becomes 0.50...1.30 g. The plausible range for the value per mass is \$79...189/g (with a midpoint of \$122/g).

The next estimate is the value per mass of gold. I can be as accurate as I want in converting from ounces to grams. But I'll be lazy and try to remember the value while including uncertainty to reflect the

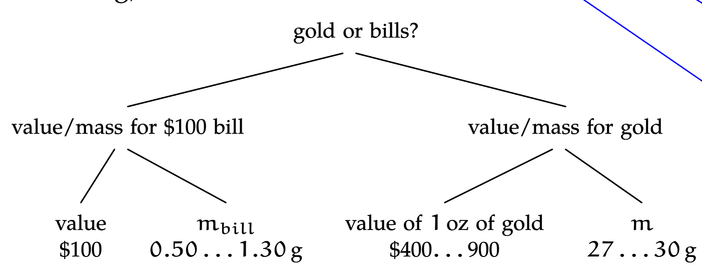

is this what we should be putting the answer boxes on our psets?

My answers are not quite so in depth, but this would be good. Alternately, you can outline your logic or assumptions or write your explanation as if you were giving hints to do the problem.

My answers are rarely this in depth - but I think he suggested only a few sentences for that purpose.

I like how this is laid out though. If gives a little more insight into how all the values came about.

This seems a lot like Sanjoy's explanations on the pset answers. I don't think ours need to be so in depth, but I guess it would be really useful for him to read exactly how our minds work!

From the grading policy, Sanjoy said, "For a reasonable effort on the homeworks, give a coherent one- or two-sentence explanation or reasoning for each problem... Write down what you could tell your earlier self in order to make finding the solution smooth sailing."

I think these long explanations are just to make sure anyone who was confused can understand it.

I think it would have been really helpful to have more explanations of this back at the beginning of the course when we were first learning how to get a feel for ranges using our gut.

We already established that it floats in water, so why couldn't we minimize this to 0.8...1g/cm^3?

Because it is so thin, my thought was that it might float due to surface tension rather than density. Like some bugs...denser than water but still float because of how their feet are arranged

wood floats, but does money?

I feel more comfortable dealing with this range, than using your original estimate of rho=1.

I say that a bill would float, so wouldn't that cut off the upper part?

and propagate them upward – thereby obtaining plausible ranges for the value per mass of bills and gold.

Problem 8.2 Your plausible ranges
What are your plausible ranges for the five leaf quantities t_{ream}, N_{ream}, w, h, and ρ? Propagate them upward to get plausible ranges for the interior nodes including for the root node m.

Here are my ranges along with a few notes on how I estimated a few of them:

1. thickness of a ream, t_{ream}: $4\ldots6$ cm.
2. number of sheets in a ream, N_{ream}: 500. I'm almost certain that I remember this value correctly, but to be certain I confirmed it by looking at a label on a fresh ream.
3. width of a bill, w: $10\ldots20$ cm. A reasonable length estimate seemed to be 6 in but I could give or take a couple inches. In metric units, $4\ldots8$ in becomes (roughly) $10\ldots20$ cm.
4. height of a bill, h: $5\ldots7$ cm.
5. density of a bill, ρ: $0.8\ldots1.2\,\mathrm{g\,cm^{-3}}$. The argument for $\rho = 1\,\mathrm{g\,cm^{-3}}$ – that a bill is made from paper and paper is made from wood – seems reasonable. However, the many steps required to process wood into paper may reduce or increase the density slightly.

Now propagate these ranges upward. The plausible range for the thickness t becomes $0.8\ldots1.2\cdot10^{-2}$ cm. The plausible range for the volume V becomes $0.53\ldots1.27\,\mathrm{cm^3}$. The plausible range for the mass m becomes $0.50\ldots1.30$ g. The plausible range for the value per mass is $\$79\ldots189/\mathrm{g}$ (with a midpoint of \$122/g).

The next estimate is the value per mass of gold. I can be as accurate as I want in converting from ounces to grams. But I'll be lazy and try to remember the value while including uncertainty to reflect the

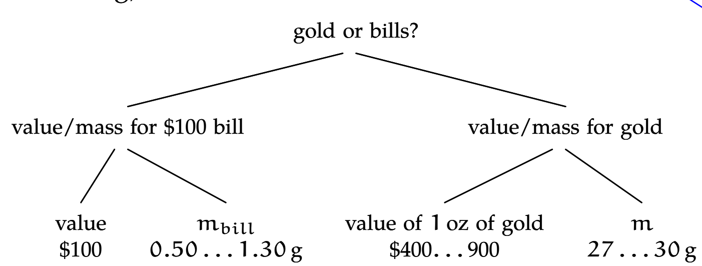

Maybe you could be slightly more explicit since this is a new paragraph and say upward through the branches in the tree.

This is an important concept that I think could have used more explanation. We went over it in class a bit and it seems crucial to getting a final range but the details were not clearly mapped out.

I'm still a little confused on the details actually...

Agreed. It's a neat concept, but the idea or purpose of this technique isn't really understood until after the example is finished. It would've been nice to explain the concept first in theory before doing an example on it.

I'm still not sure how you get these ranges? Did you just do out all the math of the previous section, or is there some shortcut?

for this propogation are you doing the same squaring/ square rooting as before?

Do you mean adding in quadrature? It's interesting to note that you can't always add errors that way, eg. if the variable in question does not go linearly or goes as a logarithm.

I would like this paragraph more if it were written with the math drawn out ... I feel like it would actually reinforce the lesson better ... rather than being an exercise in page flipping (i may be biased by the fact that going back here is _a lot_ harder than flipping pages) ... but it would be a lot better to see the numbers drawn out.

How are these endpoints arrived at? They are not merely the upper and lower bounds of the density and volume multiplied respectively.

I think he just gave a range of values that seemed likely to him. These values are much cleaner than if he merely used lower/upper bounds.

I feel like this is what we were doing the 1st couple weeks of class, just adding a probablistic element to it to make it more accurate.

How do we get these numbers from the \$100 value and the mass range of 0.5g ... 1.3g? The values of \$79 per gram and \$189 per gram are not (100/1.3) and (100/0.5). I'm confused here.

and propagate them upward – thereby obtaining plausible ranges for the value per mass of bills and gold.

Problem 8.2 Your plausible ranges
What are your plausible ranges for the five leaf quantities t_{ream}, N_{ream}, w, h, and ρ? Propagate them upward to get plausible ranges for the interior nodes including for the root node m.

Here are my ranges along with a few notes on how I estimated a few of them:

1. thickness of a ream, t_{ream}: 4 . . . 6 cm.
2. number of sheets in a ream, N_{ream}: 500. I'm almost certain that I remember this value correctly, but to be certain I confirmed it by looking at a label on a fresh ream.
3. width of a bill, w: 10 . . . 20 cm. A reasonable length estimate seemed to be 6 in but I could give or take a couple inches. In metric units, 4 . . . 8 in becomes (roughly) 10 . . . 20 cm.
4. height of a bill, h: 5 . . . 7 cm.
5. density of a bill, ρ: $0.8 \ldots 1.2\,\text{g}\,\text{cm}^{-3}$. The argument for $\rho = 1\,\text{g}\,\text{cm}^{-3}$ – that a bill is made from paper and paper is made from wood – seems reasonable. However, the many steps required to process wood into paper may reduce or increase the density slightly.

Now propagate these ranges upward. The plausible range for the thickness t becomes $0.8 \ldots 1.2 \cdot 10^{-2}$ cm. The plausible range for the volume V becomes $0.53 \ldots 1.27\,\text{cm}^3$. The plausible range for the mass m becomes 0.50 . . . 1.30 g. The plausible range for the value per mass is \$79. . . 189/g (with a midpoint of \$122/g).

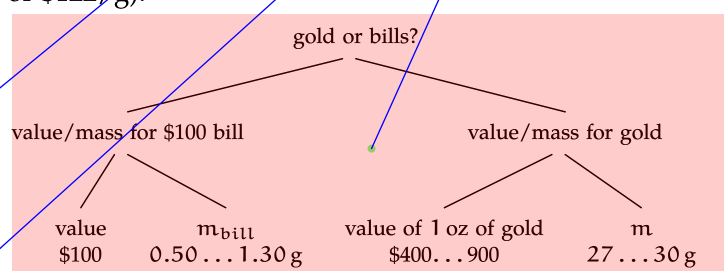

The next estimate is the value per mass of gold. I can be as accurate as I want in converting from ounces to grams. But I'll be lazy and try to remember the value while including uncertainty to reflect the

Even now, at the end of the class, I think these trees will be the most tangible thing I've learned in this class. They lay everything out so simply, a true epitome of the organization of the problem at hand.

They're a really nice formalism for something that most people try to do messily in their heads in all the time.

Yeah, the trees are especially helpful for me because when I try to work out problems in my head I tend to forget one important facet of the problem.. with the trees I can visualize the important pieces of information before I begin my estimations and calculations.

You say you'll be lazy instead and just try to do it from memory? What would be the other option? (I.e., would you just use more accurate estimation methods to convert it?)

Haha I just laughed a little. I like your bits of humor everywhere.

fallibility of memory; let's say that $1oz = 27\ldots30\,g$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400...900. The mass and value ranges combine to give \$14...32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

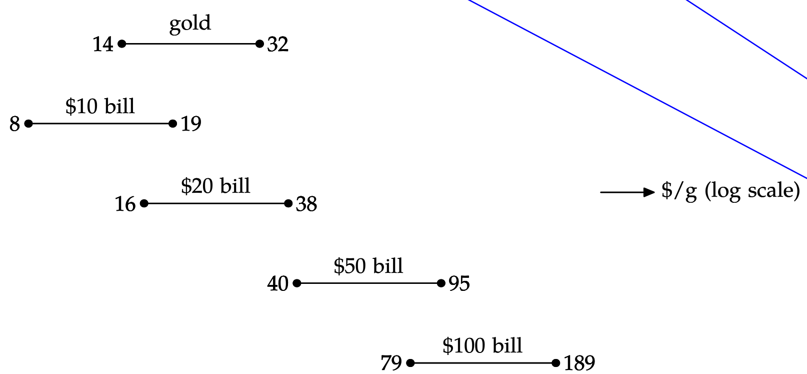

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

Comments on page 5

I am not sure if I understand the connection between this and probability? This memo seems more like a review of divide and conquer and some other methods we learned in the past.

> It uses probability to analyze divide and conquer. Perhaps I should make this example, without rangs, the vey first example in the book. Then return to it to illustrate ranges.

my understanding of everything so far is that the point of probability is to calculate the error ranges.

the placement or wording of this phrase sounds weird to me. But I understand the point you are making with it.

You mean in terms of guessing, right?

that is a very large range

A large range would reflect how certain I am about my estimate of the price of gold (not certain at all).

fallibility of memory; let's say that $1oz = 27 \ldots 30\,g$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400...900. The mass and value ranges combine to give \$14...32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

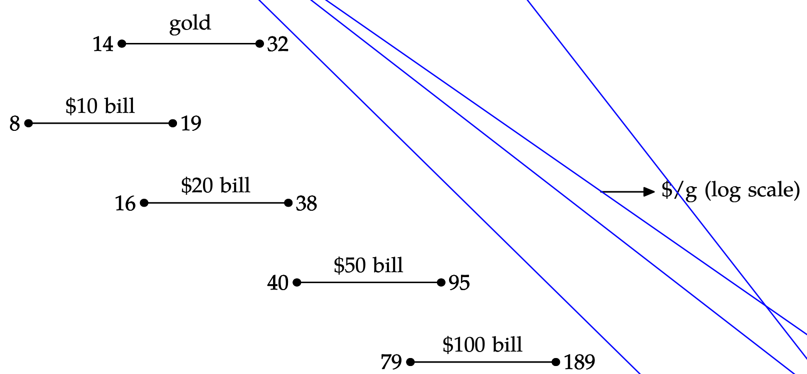

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

This still misses the actual price by a decent margin.

It's interesting that Sanjoy's lower bound is 400$ lower than his original estimate, but his upper bound is only 100$ higher, especially given that the actual number is 400$ higher than the original. Specifically, it's interesting that he was more sure it'd be lower than his estimate than higher, even given the current state of the economy.

even though he misses the actual price of gold, the guess is still close enough to tell us that if we're robbing a bank we should grab all the 100's (and probably the 50's). ie even if the price of gold was higher than what we used here, the value per mass of the 100's is still definitely larger than that of gold.

There are two answers to that. First, it's good if some of the ranges don't include the actual value. Otherwise I'd wonder if my ranges were too wide. I do want to be surprised once in a while (roughly one-third of the time).

Second, I wrote the main text in 2008 before the financial collapse. And financial collapses usually increase the value of precious metals. So if I had included that idea when I estimated the price, I would have used a higher range – maybe raising all the prices by 20 or 30%.

I just looked up the average price of gold in 2008 and it was \$871.

I would have expected this too be much higher.

I would have expected this too be much higher considering how much we value gold.

It seems to me that probabilistic divide and conquer in this case simply adds more of a range for uncertainty, a little bit more robust way of producing our ranges on our answers to hw. Graphic below really helps.

interesting visual representation of the situation.

fallibility of memory; let's say that $1oz = 27\ldots30\,\text{g}$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400...900. The mass and value ranges combine to give \$14...32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

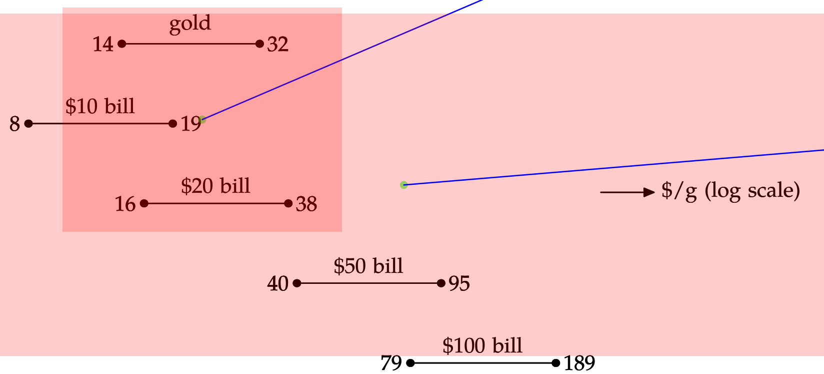

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

Earlier you say to take the gold before the 20s. This seems inconsistent.

That was when we made just one estimation. Wit this range it seems "likely" that the gold comes before the 20's.

His point here is that with just a single data point, he estimated that you should take the gold before the 20s. With a more careful calculation figuring in his relative uncertainties he realizes he's not actually so sure that taking the gold is better than taking the 20s.

I really like the picture it just might be easier to see the meaning right away if there was some kind of axes framing the data, just an x-axis.

i really like this figure

fallibility of memory; let's say that $1oz = 27\ldots30\,g$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400...900. The mass and value ranges combine to give \$14...32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

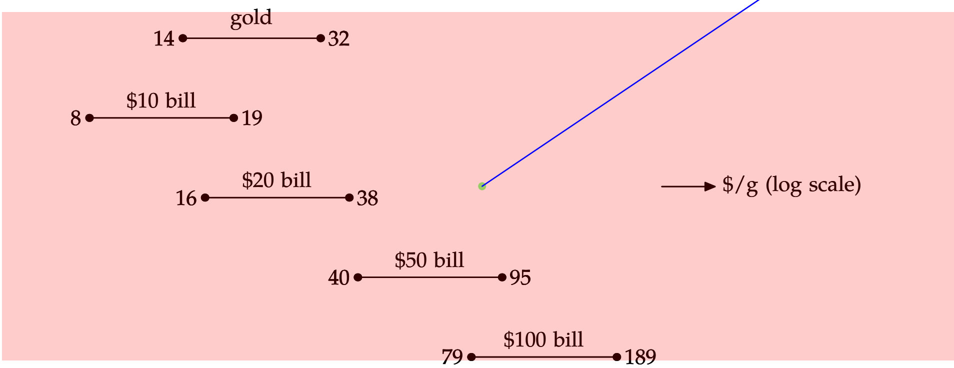

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

is this the new method we are learning in this chapter? Everything else seemed like stuff we learned a while ago

From class it seems that this chapter is focusing on ranges and creating an accurate range in order to determine the best answer. This chart is very easy to read and makes comparison much easier.

I really like this chart. It very exactly answers the question we constantly ask about how much our wrong approximations could affect the final answer.

I like the chart as well. Graphical demonstration of exactly what we just went through!

This is a beautiful graphic. It really helps to show what you were doing previously, my one question would be why you have said earlier to take the gold as opposed to the 20's? Looking at this, I would take the 20's.

I absolutely agree, the chart enables me to conceptualize the ranges and how they would effect my approximation...which through the +'s and -'s of each estimation, is usually amazingly close!

Also, will we ever quantify uncertainty in other ways than ranges?

I agree, this diagram really helps illustrate the uncertainty

I really appreciate the chart as well. It makes more sense than explaning it in words and is pictured better as well.

I agree that this diagram is really helpful, but it misses one point, which is that the value/mass for each of the bills is not independent – they're all the same size and (roughly) the same material, so knowing where in the range the \$100 bill lies will tell you where the \$50 and \$20 and so on lie. This chart, however, makes it seem like there is some change that a \$100 bill is worth less than a \$50, which is obviously false.

fallibility of memory; let's say that $1oz = 27\ldots30\,\mathrm{g}$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400...900. The mass and value ranges combine to give \$14...32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

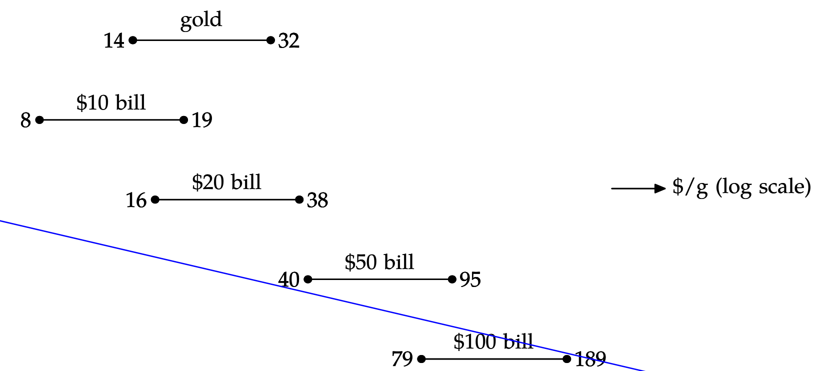

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

That's a good point. It's hard to display joint probability distributions, and this diagram "solves" the problem by ignoring all the correlations. But it should be sufficient information for the original question about what to take when you are in the bank vault.

The correlations are useful once you find out more information. For example, suppose that you later find out the true mass of a bill. That would shift all the ranges for the bills. But that information would probably come only after you leave the bank vault (unless you bring a scale with you!), so the correlations wouldn't affect your decision about what to take.

I really like this picture. It sums up the entire problem in an easy to understand way.

I think what could make this picture a little better woud be to put what we got and then the errors on both sides. I think it's called a candlestick or something in that respect, but otherwise, a very helpful picture, and I wish we had had mroe of these.

yep I agree, it is a great way to visualize the ranges.

This is a good, clear way to show us how you analyzed the results you got.

why did you decide to put this on a log scale?

I think before calculating, I kind of guessed this answer, but it was cool to follow through the 6.055 method.

fallibility of memory; let's say that $1oz = 27\ldots30\,g$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400...900. The mass and value ranges combine to give \$14...32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

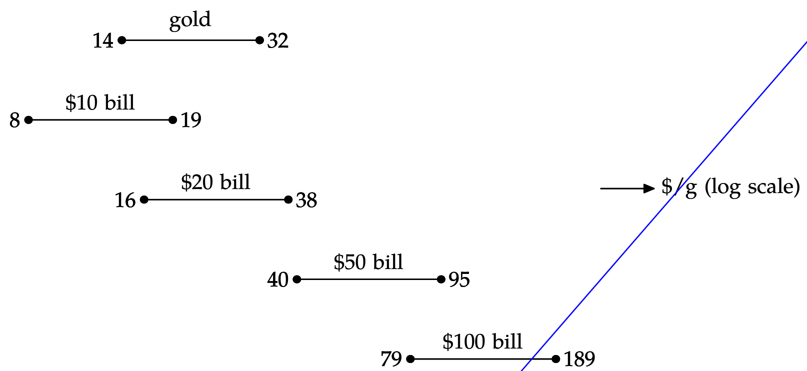

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

I guess I was wrong. I'll remember to take the bills instead of the gold next time I rob a bank.

Huh, interesting, never would have expected the difference (even accounting for uncertainty) to be so big. Value of gold per gram is roughly the same as that of a \$20 bill? No wonder the price of gold has been going up...

I guess banks should carry mainly gold bricks as opposed to bills if they want to lose less money in the event of a robbery.

there are, however, the biggest thing that this doesn't take into account is that banks dont keep gold bricks. they have safety deposit boxes – full of jewelry (and other things they're less likely to report missing) ... this means that (1) they don't necessarily know the "size of your score" and (2) what you're stealing is not generally just gold...it includes all of those gems that are set in that gold ... which is generally worth _much_ more than the gold its self.

I like how we're using probability to double-check our approximations.

This use of probability really makes me more confident in using "everyday values" for divide-and-conquer methods.

I don't feel like we're using probability though. These are just ranges, not probabilities. They are somewhat similar, but insofar as the calculations go, no explicit probabilities were used.

So by "using probability", does that just mean establishing a range of values that you're reasonably sure are correct (like 2/3 probability you are correct, and 1/3 probability you are incorrect)? Personally, that doesn't feel like applying probability. That just feels like the basic estimating skills we were taught from the beginning.

this is an awesome problem- I really really liked it

this section could theoretically be right after divide and conquer and allow for you to calculate errors for al of the chapters afterward

fallibility of memory; let's say that $1oz = 27\ldots30\,\text{g}$. This range spans only a factor of 1.1, but the value of an ounce of gold will have a wider plausible range (except for those who often deal with financial markets). My range is \$400…900. The mass and value ranges combine to give \$14…32/g as the range for gold.

Here is a picture comparing the range for gold with the ranges for US currency denominations:

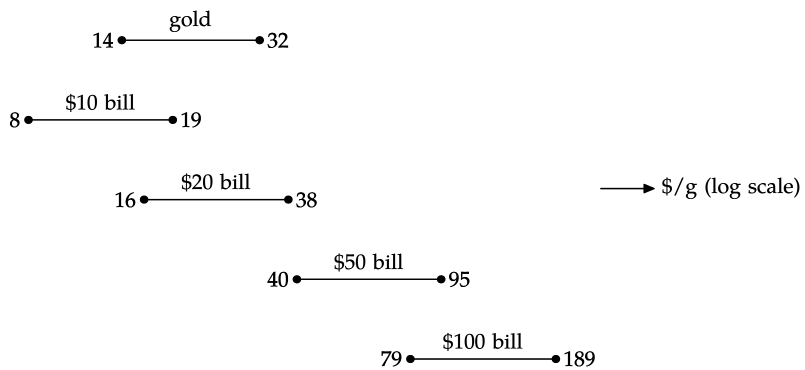

Looking at the locations of these ranges and overlaps among them, I am confident that the \$100 bills are worth more (per mass) than gold. I am reasonably confident that \$50 bills are worth more than gold, undecided about \$20 bills, and reasonably confident that \$10 bills are worth less than gold.

I really appreciate this example as a stepping-stone from the previous reading. I feel like it added just enough additional complexity so it didn't feel like we were overworking the problem, while at the same time wasn't very difficult following the simple example from the previous reading.

I also thought that this reading was excellent. I really enjoyed it, and it was very well-explained and tied the probabalistic methods and divide-and-conquer methods together. Well done.

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

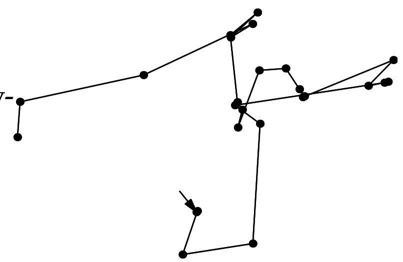

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

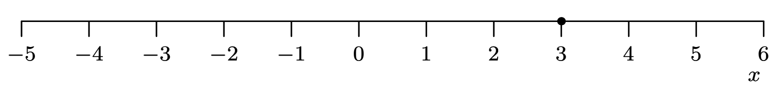

This was a rather lengthy reading, at least compared to the other ones. I feel like you tried to give me too much information at one time.

> That's probably true. I'll aim for better planning next time...

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

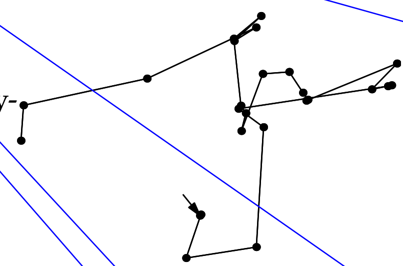

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

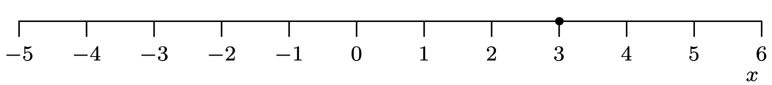

Read this section for Friday (memo due Friday at 9am).

maybe it's better to give some context because not everyone knows about the game

I think the rest of the examples convey the idea of something moving through a field toward an eventual destination.

I agree. It's okay that this reference might not reach everyone because there are so many others. But I do think this example is pretty well known.

what is a random walk? maybe your going to define it later

that's all i did in 4th grade

that was what we played after we were "too big" for train.

What is train?

War!!! Yes!!! Wonderful.

I though train was still more fun. War got boring =D

I think it would be more effective to remove this sentence, and then just start with the questions. The questions make an easy transition to the next paragraph.

I agree, I was in the mindset of 6.00 graphs and nodes. And then I got confused for a bit.

Yeah, I too agree. This sentence seems somewhat out of place compared the rest of the content.

I have been introduced to this concept before in my probability class but I am really interested in learning it in the context of 6.055.

I never thought about that...

huh? the walk or the card game? this probs should get reworded.

I think it's pretty clear it's referring to the card game.

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

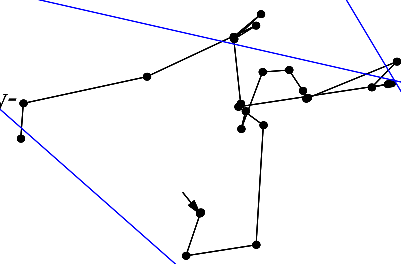

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

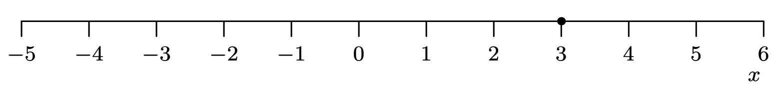

It would work better if this were another question the way the War example is.

i agree. this is a strange transition between thoughts.

Really random transition! But I am excited that we might talk about neurons and signal transduction

FOREVER and I felt like I always lost and for someone who really likes to win that brings back some bad memories

This paragraph seems very disjointed. You immediately go from talking about random walks to card games to neurotransmitters. The switchover to neurotransmitters is the worst. The sentence looks like you've been talking about neurotransmitters for a paragraph, when you really just started it. You should start this paragraph with the first sentence. Then define a random walk. Then say that some examples are card game length, neurotransmitters, etc.

yeah I agree...this list of unrelated examples at the beginning without a definition of random walks is pretty confusing.

Be careful here. You're implying that the synapse is at the neuromuscular junction. If it is, there is no "random walk"–it is guaranteed that your muscle will twitch. You don't want to rely on probability when a lion is chasing you, etc.

I feel like the bigger problem here is that he left out all the stuff in between. He goes very detailed about the neurotransmitter, then jumps to a leg twitching. Why not just say "and a signal is passed to the next neuron" or something more accurate.

more than you wanted to know about vesicles: http://en.wikipedia.org/wiki/Vesicle_%28Biology_and

Every time I've learned about random walks they mention the "drunken sailor." Perhaps it should be included with these examples?

I agree with this - its always presented as a man walking left or right with the question of will he get home. A quick mention of this could jog some people's memories

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

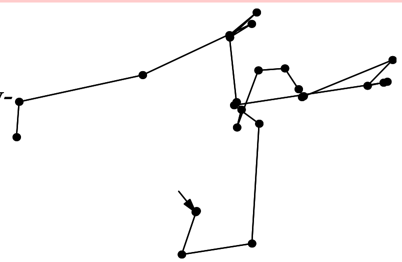

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

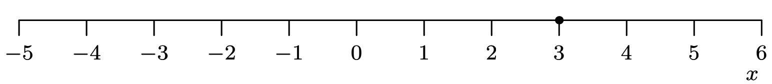

Um... what?? This paragraph is worded VERY strangely, with no transitions whatsoever between thoughts. I don't know what the exact definition of a random walk is, and this opening paragraph just confused me. It would be better to put the formal definition first.

I agree

I actually liked this paragraph. It gave a lot of examples demonstrating that random walks are everywhere, giving me some insight into what a random walk may be, before telling me the definition.

I agree with the original comment, that the paragraph was worded strangely. Perhaps more obvious transitions between examples would have helped the flow and help us understand what a random walk is before the formal definition is given.

I agree with the second comment. I liked the paragraph and although it did explicitly define what a random walk was, I was able to kinda get the idea from the multiple examples given. Plus I kinda assumed that a random walk was kinda self explanatory: the path that results from someone or something walking randomly with no logic....completely arbitrary. The wording doesn't bother me because it sounds very conversational

Yeah I would rather be given something to think about then be told a formal definition. It helps me make the idea my own and really understand it.

I like the idea of this paragraph as an opener, but I agree it's really choppy and kind of hard to follow.

Talking about random walks reminds me of the example from the beginning of the semester when we questioned whether a person climbing up and down a mountain would ever be at the same place at the same time, I really liked that example so I am hoping this discussion is just as interesting!

This would be a good example to mention again in the intro paragraph to kind of tie this section w/ examples from the past.

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

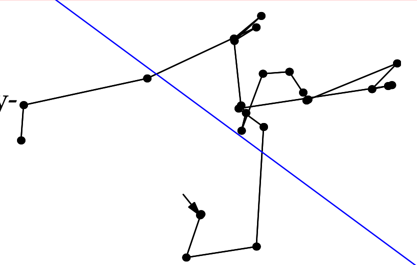

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

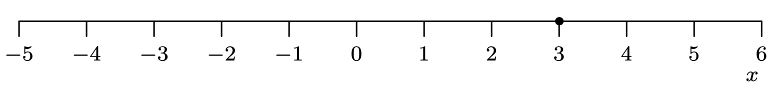

I don't see how this example has to do with random walks...

Yeah I definitely did a double-take...does this actually relate?

Just guessing...the vesicle bobs around at the terminal bouton before releasing so how fast it releases is (to some extent) dependent on random walk (how long it takes to make it to the membrane)...cold would be random motion of colder/slower particles and how often they contact your skin and slow down skin molecules...

but the vesicle is directed by receptors on the post-synaptic end...so its not really random movement

This box was around the "winter day" example, I believe. In this case, my guess is that the random walk determines how long it takes for a cold air molecule to get to your skin. (Well, perhaps how long it'd take molecules to hit a large enough area of skin to make you feel "cold".)

Referring back to the neurotransmitter, the molecule is directed but that doesn't mean that it has to necessarily make the same path from point A to B everytime. It just has to get to B.

Maybe I don't understand "random walks" yet, but I don't think this one fits in. It doesn't seem random at all. Every time you don't wear enough clothing, you'll always feel cold.

I remember the first year I was at MIT. Every time my friend walked outside during the winter he would say "aww it's not that bad" we'd hardly walk 50 yards before he'd be freezing.

I think it's referring to the motion of molecules, since all temperature is, is how fast the molecules are moving. Since it's cold they aren't moving as fast and are probably easier to approximate a random walk for.

That's the beauty of randomness: that it can be understood. The molecules move randomly, but there are so many of them that you can predict what they do in the aggregate – and they always make you feel cold.

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

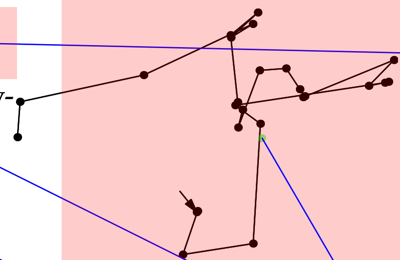

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

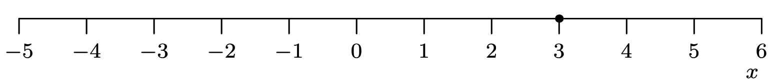

Does War fall into this category? It makes the point in a way that makes sense and is easy to relate to, but I"m not sure I think of War as a physical situation the same way as I do the movement of molecules

I think the comparison is that in War, every play is random. You don't know if person A or person B will win.

And the physical situation is the actual cards themselves moving from one player's hands to the other's.

After giving examples, perhaps it would be a good idea to state a formal definition of random walks. I've never heard of it before and would want to make sure that I am coming to the right conclusions.

That might be a good idea. Although formal definitions are, I've realized, alien to my way of thinking. I'm not sure what the formal definition of a random walk is, but I "know it when I see it." The picture I use in my own head is, "it's what a gas molecule does as it zooms around, bouncing off other gas molecules." The main point is that the motion after a collision is in a completely random direction (i.e. no information about the previous path remains).

These physical situations are examples of random walks, but I only have a vague, but not definite idea, of what a random walk is. I think it would be very helpful to be given a definition.

aw, your e-mail made me hope that you'd be using an example of an aimless stroll.

i'm just now figuring out that that's not what you meant...

This is an effective diagram documenting randomness.

I actually think this is somewhat unnecessary, though I guess it serves it's purpose.

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

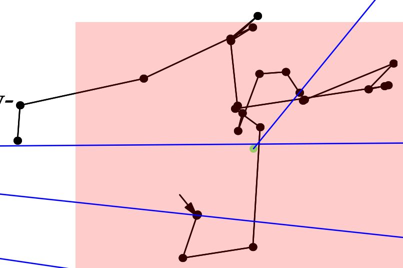

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

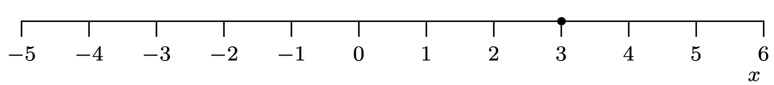

are you going to reference this diagram? it's different than the "regular walk" described below.

Agreed, it seemed a bit "random" without any sort of explanation/ reference to it.

I think it's okay if you don't reference it. It adds to the paper and it is pretty clear what this figure suggests.

Haha maybe it's a random figure showing a random walk. Clever

Regardless of the fact that it seems loosely related to the subject at hand, I still think some explanation would be helpful.

When I think of a random walk, I imagine motion that is undirected. So I was initially confused by the idea that war is random walking. But, the cards are played against one another in a somewhat random order.

Once again, I don't really get the random walks. I feel like random walks should not behave in a pattern that we can figure out.

A definition would be good in here.

I agree. I can get an idea of what a random walk is from your examples, but a formal definition is the only way outside of a lot of examples or reading through the whole section that I could really be sure of the definition and I think it would be better to know what you're talking about before reading the section.

don't they behave "randomly"?

I think their behavior can be predicted (ie brownian motion).

Again with the cannot uncheck a question thingie! Yah, i think he means how they behave so we can model it. "random" isnt exactly a model

is there a diffusion equation for the game of war?

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

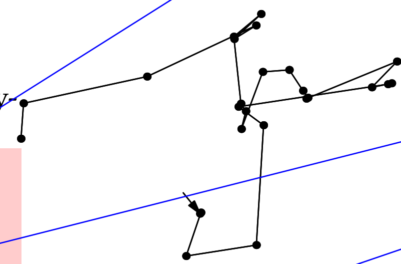

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

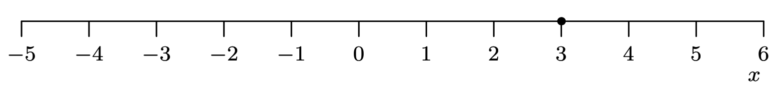

I like this description of the process we're going to take.

I agree. For me, it's really useful to have our plan-of-attack spelled out beforehand, so that we can better understand what we're supposed to be taking away from each part of the problem.

I've been learning from the previous reading memos! One of the main themes is that roadmaps are greatly valued, so I've been trying to add them to subsequent readings.

Of all the examples, this is the only one that I know what we are going figure out from this randomness

Cool! I was looking forward to this when we talked about it in one of the first lectures.

Sounds like it.

Wait why is it so complicated?

Because they can move any direction for any length they want. There's no real way to easily predict where they'll end up without using any tricks.

It's complicated in terms of evaluating the behavior. The elements are too variable, making the situation very complicated.

I feel like I haven't gotten an explicit definition of a random walk, so I'm a little confused (particularly because the situation being complex makes me wonder if there are rules or something)

The paragraph right after your highlighted box starts defining what a random walk is. Although, it doesn't explicitly make that clear.

????

This sentence seeems somewhat awkwardly worded. I'm not really sure what it means here.

yea i just don't get it

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

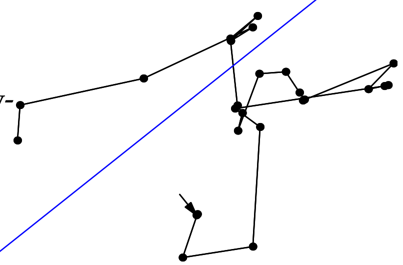

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

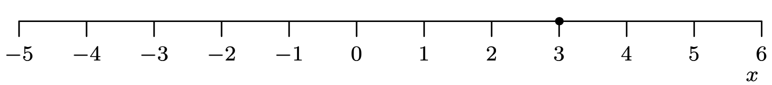

The first example of a random walk was the game war. I thought that was a great example of what we are trying to model. This first part of the section seemed a little hard to follow as I had the idea of 'war' stuck in my head. Would it be easier to understand for the reader to model the example on 'war'?

Agreed. Also, assuming that most us have played war before, it would be easier to picture the cases.

I've been thinking about "war" a bit more. The randomness is that each person turns up a card, and you don't know whose will show the bigger number. But, as you win more cards, you are likely to be getting smaller into your pack, so your probability of winning depends on how many cards you have won. Thus, history matters – which is not how a true random walk works. So "war" may not be a 100% valid example. More thinking required...

or that they have a fixed random chance of turning in a new direction after a certain distance.

I agree with your explanation. Because he's only using collision to refer to the particles example named earlier.

Yeah, this makes a lot more sense - like a person walking down the street and deciding to continue or turn around at each intersection.

I'd like an explanation for why we made the assumptions we did.

I am a little confused on what were talking about here? what is this random walk?

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

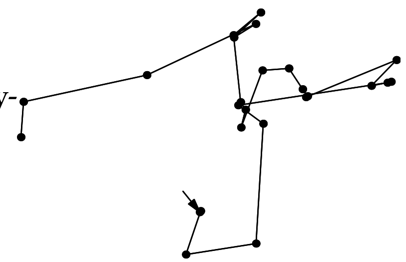

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

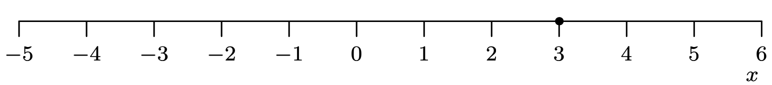

This seems like it is making the movement much less random..?

I think it is keeping things random but taking some of the generality out of it. Basically, look at a very simple easy to analyze case of a random walk. (it is indeed possible for a walk to be random and follow all of these assumption, although its probably very unlikely). After we look at the simple case then we can try to generalize. Sort of like the solitare example from several weeks ago.

sort of, but we can always build off of this. Say we solve it for one axis, then 2, then 3... we could begin to extrapolate to infinite dimensions, and this would again return us to any direction is possible. We could also do something similar for distance. but we have to find somewhere to start.

I agree this is a little bit confusing but I think I'm starting to understand what he's getting at.

This is confusing. Is random one-dimensional motion moving back and forth on a line?

So easy cases and lumping! I ilke how the methods were addressed.

only by one unit (I believe the following implies)

Thank you for the clarification - I was a little confused what our fixed distance was.

any particular readon its at 3? are we assuming is been walking for a while?

I don't think there's a reason, its random. He might have been going for a while and ended up at 3 or he might have happened to go right 6 times in a row as soon as it started.

I think it's just clarifying the starting condition.

I just picked that spot randomly, but I should explain and use that information in the paragraph.

8.3 Random walks

Random walks are everywhere. Do you remember the card game War? How long does it last, on average? A molecule of neurotransmitter is released from a vesicle. Eventually it binds to the synapse; then your leg twitches. How long does the molecule take to arrive? On a winter day, you stand outside wearing only a thin layer of clothing. Why do you feel cold?

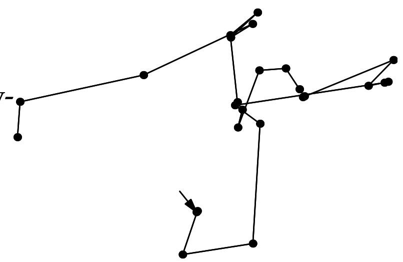

These physical situations are examples of random walks – for example, a gas molecule moving and colliding. The analysis in this section is in three parts. First, we figure out how random walks behave. Then we use that knowledge to derive the diffusion equation, which is a reusable idea (an abstraction). Finally, we apply the diffusion equation to heat flows (keeping warm on a cold day).

8.3.1 Behavior of regular walks

In a general random walk, the walker can move a variable distance and in any direction. This general situation is complicated. Fortunately, the essential features of the random walk do not depend on these complicated details. Let's simplify. The complexity arises from the generality – namely, because the direction and the distance between collisions are continuous. To simplify, lump the possible distances: Assume that the particle can travel only a fixed distance between collisions. In addition, lump the possible directions: Assume that the particle can travel only along coordinate axes. Further specialize by analyzing the special case of one-dimensional motion before going to the more general cases of two- and three-dimensional motion.

In this lumped one-dimensional model, a particle starts at the origin and moves along a line. At each tick it moves left or right with probability 1/2 in each direction. Here it is at $x = 3$:

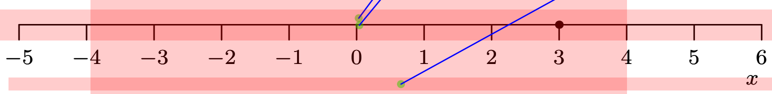

So is this number line a simplification of the random walk? why did we limit the distance it can travel and then say it can only move in one dimension? I don't really see where we're going here

> This is a special case of the general random walk (where the particle can move a random distance and in a random direction). Here, it's been restricted to move a fixed distance (a unit distance) and in one of only two directions (left or right). The analysis is simpler but the main results transfer to the more general case.

It may be helpful if you had several of these diagrams...this will simulate the movement. perhaps a actual simulation

What am I supposed to be getting out of this diagram? Its just a number line with a dot at 3. I'm usually all for diagrams, but this one wasn't particularly helpful

It's just so you can visualize moving left or right at every step.

The diagram explains the above situation. You start at x = 0. Then, after the first time step, you're at x = 1 with probability 1/2 or at x = -1 with probability 1/2. That is, at each time step, you move to the left or to the right by one with probability 1/2 each. The diagram is supposed to help you graphically understand the random walk described above.

specifically it is the set of positions that a particle moving in a one-dimensional random walk can follow. our model assumes discrete, integer steps. also, no hops and each time interval corresponds to one step with equal probabilities in available directions.

That's exactly what it's supposed to be! This just lets you visualize two directions, left and right with equal proability.

Does this imply that the model is discrete?

> Right, after using lumping, the continuous model turns into a discrete model.

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

Typo. Should read: "The expected position is the average..."

I vaguely remember this notation, but you might want to point it out a little more strongly...that one means average versus actual

I get what this word means in this context but I think it could be potentially confusing if one didn't know the concept.

But isn't it clarified by the aside after the em dash?

what do you mean by this? that the expected position is the origin? or that it's the same for each walk?

That the expected position is to return to baseline (which in this case is not the origin, its at 3)

Stating this more explicitly as in the comment at 1:25 would be helpful.

He means that the expected position is the same for each walk for which you start at the same point. The starting point (also referred to the baseline in the above comment) for this problem was the origin, x = 0. So in this problem, the expected position is in fact the origin. The above comment at 1:25 is incorrect by saying the baseline is x = 3, because that is not what the problem states above. It simply shows the particle at x = 3 after a certain amount of ticks.

That makes much more sense than the previous comments.

What about the variance?

You might want to just say expected value of the position or explain what the first moment is.

yea what is the "moment?"

I think he means that since the average/expected value is constant, you can just throw out the 'n', since it doesn't chance. So instead of <xn> you get just <x>.. By "first moment" i think he just means where it starts, since that is where the average will lie.

This might help: http://en.wikipedia.org/wiki/Moment_%28mathematics%29

Different moments provide different characteristics of a population. The first moment is the population mean.

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x^2_{n+1} \rangle = \langle x^2_n \rangle + 1.$$

In other words,

$$\langle x^2_n \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x^2_n \rangle \propto t.$$

is that what the angled brackets mean?

I don't believe so. I believe the angled brackets are simply used to distinguish the action position from the expected position.

Also, why is it an invariant? I thought it was random...?

I think he's talking about the expected value of x i.e. <x> and that is zero and since it is zero (constant) it is invariant

It's good to see invariants making a comeback here.

But I don't understand... what's the first moment? I get that its invariant but I'm lost from there.

Ah. I have just been informed that the first moment is simply the expected value of x, so just the average. If this is the case, just take a couple words to mention that terminology.

I agree. Introducing the term "first moment of the position" is kind of unnecessary (since we already have several other terms for this quantity) and vaguely confusing. If it's really important that it be kept, it would be helpful to define it a little more explicitly.

Moments are a specific definition in probability. Is it worth specifying what a moment means here?

He means moment as in the next physical movement in time right, but that doesn't seem right, what is it?

Ya, this term is ambiguous to me in this context. I think you should clarify what you mean.

I should do that. The first moment is the expected value of x^1. The second moment is the expected value of x^2, etc.

I don't understand why the 2 here is a superscript (as if to indicate a squaring) rather than a subscript (which would indicate a second walk)?

because it is squared.

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

Why is it squared?

Should that be a subscript ?

It shouldn't be a subscript but I'm still confused why we are looking at squares?

Why are we interested in expected square position- are we trying to calculate variance and expected value?

I think its because it has already moved once, <X> and now it is moving again, the same distance but we don't know the direction, so distance <X^2>

The second moment is equal to the variance, which I think equals the position squared.

The second moment will give us different information than the first moment, and in this case more useful information, since we already knew the information provided by the first moment. And yes, the second moment is equal to the variance which is equal to the average value of position squared, as the notation suggests.

What does he mean by second moment? I see us simply moving in increments of one. It doesn't make sense to me that we use x^2. What is the reason to calculate variance?

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

the n-th moment is defined as:

<x^n>

or "the expected value of the nth power of the variable"

like people said previously, the 1st moment happens to be the mean, the 2nd moment happens to be the variance.

we are calculating mean and variance because it tells us information about the nature of a random walk. Since there is equal probability of going in any direction, the mean is wherever the particle starts at. The variance, on the other hand, is the square of the standard deviation, which tells us that 67% of the time, the particle will fall within one standard deviation after moving for a time t. From the formula, we see that one standard deviation is the square root of t.

The primary reason we use x^2, as you asked, is because it will take into account movement in the negative direction. Let me put it this way:

The particle is equally likely to move in either direction, so if we run this experiment for many trials, the mean, or <x>, will average out to wherever the particle started at (e.g. x=0). But to get an idea of what the average "variation" in distance from the mean is in those trials (hence the term variance), we need to take how far the particle has gotten from the mean in each trial and average all those distances. Trouble is, the particle is equally likely to have negative distances as positive distances, and with equal probability of going in either direction, such an average would come out to zero. Surely, you wouldn't believe taht the average distance traveled in each trial is zero. Thus, mathematically, what we do is square the distance traveled in each trial to get rid of negative distances. Then we average that to get the average squared distance traveled. Finally, all we have to do is take the square root to get a more practical interpretation: the standard deviation, which tells us about how far each particle travels in time t from the mean.

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

I think the important part here is linking what we're doing to what people reading will understand more intuitively. Saying that the square root of <x^2>=stdev would help a lot since most people reading this would understand standard deviations.

We are looking for the expected value of x^2, i.e. <x^2> (not the square of <x>). Since x^2 is always nonnegative, it's expected value is also nonnegative (and is in fact positive for n>0).

To find <x^2> by brute force, we'd need to find the probability of x^2 being 0, of it being 1, of it being 2, etc. and weight 0, 1, and 2, etc. by those respective probabilities. I should definitely give more details in the book, and then say, "But we're going to do it by easy cases and guessing instead of by brute force." But at least then everyone would know what we are trying to guess!

At this point it became pretty confusing for me. The <x^2> thing didn't really make sense when you started adding in the subscripts.

Maybe you could explain what <x^2> means in the reading.

i think adding <x^2> = 49 here would be good.

What is the difference between expected square position and squared expected position (in terms of their purpose)? I get that the former is a way of finding averages. Is the latter significant in any way?

This is actually a very subtle but really important difference. I feel like there should be a short explanation about it.

Yeah, I was a bit confused at first.

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

It is important you explain this difference because the number 1 confusion in probablity is differentiating <x>^2 from <x^2>

I'm confused what you mean by this - can you elaborate? (relatedly, I'm also confused about this portion of the reading and what the x^2 is supposed to mean.)

The difference here seems to be:

<x^2> = 50 <x>^2 = 49

This is because <x> = 7, and squaring it (or finding the squared expected position) is equal to 49.

This would be a really great additional way to write this sentence!

I'm confused...how is the number 50 useful? it just means you are in the same position 7.

why wouldn't this be 1/2(7^2)=24.5. why are we using these two quantities.

Not really sure where the half came from?

im lost on how we got here. the average of the squares?

yes, what is the basis for this operation?

Yeah, I'm very confused on this line, on the whole x^2 thing, and on the bracket <> notation. I think a review of probability formulas and terminology would be appropriate before diving into this section.

So he's trying to find the expectation of the squared position. If we're at x = 7, then we'll move to x = 6 with probability 1/2, and move to x = 8 with probability 1/2. Remember, the value x is the position. So, E[x^2] = (probability x = 6)*(6^2) + (probability x = 8)*(8^2) = 1/2*(6^2) + 1/2*(8^2) = 1/2*(6^2 + 8^2) = 1/2*(36 + 64) = 1/2*(100) = 50. The bracket <> simply denotes expected value.

I understand this, but I don't understand how multiplying the probability by the square of the next possible position is helpful t determining anything. In a one dimension scenario does that sort of regularization account for anything?

I don't understand why we have an x^2.

how does finding this help us? What does 50 have to do with anything?

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

I'm SO confused what's going on here, I feel like this section isn't as well explained as other sections with tougher information.

I agree - i'm not really sure where this is going either...

I get the concept mathematically, but not conceptually.It seems like you are trying to make a point that everything is relative. I'll keep reading.

Had the same feeling - was a bit lost but want to see where this is going...hopefully i get the destination

so what does this mean? i don't know much about probability or what we're looking for here.

I'm wondering the same.

Hm yea I forget this stuff, but I think a big reason to use squared anything is to get rid of the differences between positive and negative changes. So my guess about what this means is that the absolute value of the expected position increased by one. Or something like that maybe.

Keep reading. You'll see that he's trying to find a pattern for the invariant. We literally just wanted to find out that the expected squared position increased by 1.

I understand that the squared position increased by one. But why does this matter?

Well, if you read above, it clarifies that this is easy cases - the squared (or second moment) is the easiest to analyze "in special cases"

how is it increased by 1

You should mention about how 7ˆ2 was 49. it took me a bit to figure out what you meant by plus one

I am glad you wrote this note, because I was rather confused myself.

I'm SO confused what's going on here, I feel like this section isn't as well explained as other sections.

oops, sorry there are multiples of this comment, my firefox malfunctioned...

I'm SO confused what's going on here, I feel like this section isn't as well explained as other sections.

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

can someone please define what the expected squared value means?

It means exactly what its algebraic formalism gives it as, and it's a stepping stone to finding the variance and standard deviation of a distribution.

I just realized that you were taking the difference of X^2 and <x^2>. But I dont know whY?

I know that this is the equation for variance in probability, but I don't know why either. I think this is a different way of proving that variance in the difference between these two magnitudes, but I'm still confused by his method.

that's actually really cool – it makes sense in retrospect that (n+1)^2+(n-1)^2 = n^2 + 1 but i never thought about it like that

Why is the expected squared position important, why no the expected cubes position, or some other function of the expected position.

Why do we take the avg of the squares?

I feel like just using numerical examples is not convincing - what would be better is an algebraic argument. You can show x^2+1 = .5(x-1)^2+.5(x+1)^2, which might be more intuitive for MIT students.

Oh, I really like your explanation here. I didn't get why it kept increasing by one every-time, but I understand now that you say x^2+1 = .5(x-1)^2+.5(x+1)^2. I would definitely consider adding this explanation to the text.

I thought the explanation was alright, though I was lost until the equations were put down. After they were put down it was a "Oh, wow, cool" moment.

Let the position after n steps be x_n, and the expected position after n steps be $\langle x_n \rangle$. The expected is position is the average of all its possible positions, weighted by their probabilities. Because the random walk is unbiased – motion in each direction is equally likely – the expected position cannot change (that's a symmetry argument).

$$\langle x_n \rangle = \langle x_{n-1} \rangle .$$

Therefore, $\langle x \rangle$, the first moment of the position, is an invariant. Alas, it is not a fascinating invariant because it does not tell us anything that we did not already understand.

Let's try the next-most-complicated moment: the second moment $\langle x^2 \rangle$. Its analysis is easiest in special cases. Suppose that, after wandering a while, the particle has arrived at 7, i.e. $x = 7$. At the next tick it will be at either $x = 6$ or $x = 8$. Its expected squared position – *not* its squared expected position! – has become

$$\langle x^2 \rangle = \frac{1}{2}\left(6^2 + 8^2\right) = 50.$$

The expected squared position increased by 1.

Let's check this pattern with a second example. Suppose that the particle is at $x = 10$, so $\langle x^2 \rangle = 100$. After one tick, the new expected squared position is

$$\langle x^2 \rangle = \frac{1}{2}\left(9^2 + 11^2\right) = 101.$$

Yet again $\langle x^2 \rangle$ has increased by 1! Based on those two examples, the conclusion is that

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

In other words,

$$\langle x_n^2 \rangle = n.$$

Because each step takes the same time (the particle moves at constant speed),

$$\langle x_n^2 \rangle \propto t.$$

This can be just as easily be seen by thinking of x_n as the sum of n random variables y_1...y_n, each y_i=1 or -1.

Since each var(y_i)=1, var(x_n) = var(y_1+...+y_n) = n.

(And since <x_n>=0, var(x_n) = <(x_n)^2>.)

p.s., this approach would also apply to more general walks, for which we only need to know that <y_i>=0, and the variance of each y_i.

This is more applicable when the distance at each step is some continuous distribution (like in the log of our errors during estimation, for example, where each var(y_i) log(r_i)^2).

I don't make this connection. n is the number of moves?

Yes, I think so. so x(n) describes a sequence of events, and x(n) is the position at time n.But i still don't understand how this works...what is n in this case? (from the example of x=10, <x^2>=100?)

So if we start at the origin, is the <xn> –> x=0 at n=0 or n=1?

but <7^2> was 50 ....? how is that n?

I think you may be confused about the notation. n is not where you start, it's the number of tick marks that you move away from a starting position...so in the examples above, we moved once and found that the expected value increased by 1.

So simple yet hard to realize.

I agree, it is rather fascinating.

Yeah I think it's crazy once you think about it. In a good way of course.

The result that $\langle x^2 \rangle$ is proportional to time applies not only to the one-dimensional random walk. Here's an example in two dimensions. Suppose that the particle's position is $(5, 2)$, so $\langle x^2 \rangle = 29$. After one step, it has four equally likely positions:

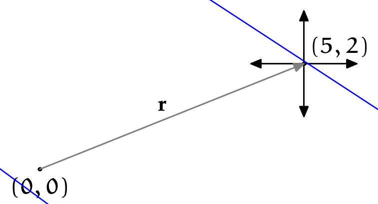

Rather than compute the new expected squared distance using all four positions, be lazy and just look at the two horizontal motions. The two possibilities are $(6, 2)$ and $(4, 2)$. The expected squared distance is

$$\langle x^2 \rangle = \frac{1}{2}(40 + 20) = 30,$$

which is one more than the previous value of $\langle x^2 \rangle$. Since nothing is special about horizontal motion compared to vertical motion – symmetry! – the same result holds for vertical motion. So, averaging over the four possible locations produces an expected squared distance of 30.

For two dimensions, the pattern is:

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

No step in the analysis depended on being in only two dimensions. In fancy words, the derivation and the result are invariant to change of dimensionality. In plain English, this result also works in three dimensions.

In a standard walk in a straight line, $\langle x \rangle \propto$ time. Note the single power of x. The interesting quantity in a regular walk is not x itself, since it can grow without limit and is not invariant, but the ratio x/t, which is invariant to changes in t. This invariant is known as the speed.

In a random walk, where $\langle x^2 \rangle \propto t$, the interesting quantity is $\langle x^2 \rangle / t$. The expected squared position is not invariant to changes in t. However, the ratio $\langle x^2 \rangle / t$ is invariant. This invariant is, except for a dimensionless constant, the *diffusion constant* and is often denoted D. It has dimensions of L^2T^{-1}.

This qualitative difference between a random and a regular walk makes intuitive sense. A random walker, for example a gas molecule or a very

will this work even when the particle moved at fractions of tick marks?

I still dont understnad why we are loking at xˆ2

It's essentially (glossing over some technicalities) the variance of your location after n steps. Variance is sort of a measure of spread, and its square root is standard deviation. Since this is a random walk, you don't know exactly where you'll end up, but the variance/standard deviation gives you an idea about how far you might be from your starting point.

Well put.

Very helpful. I could have really used such an explanation in the text somewhere.

at this point I realized that following your x's is really hard sometimes...you should use subscripts for all of them, exp when the eq for obtaining them is different .. it took me _way_ too long to figure out why this wasn't 14.5

Note, and remember for later, that x is a vector magnitude here (unless you mean r instead of x)

Maybe show these on the diagram below?

It looks as though that is what the 4 arrows are meant to show. They just don't have the coordinates associated with them.

I think he should have also mentioned he's using lumping here since really you could move in more than for directions.

has the particle always been 2 dimensional? initially i thought everything was (x,0)

For the particular example where you placed your comment, it is two dimensional. Previous examples were 1D though.

I don't really understand what is going on...this is not really making sense to me. I guess I don't understand what an expected value is.

Squared distance makes a little more sense to use in 2 dimensions than 1.

The result that $\langle x^2 \rangle$ is proportional to time applies not only to the one-dimensional random walk. Here's an example in two dimensions. Suppose that the particle's position is $(5, 2)$, so $\langle x^2 \rangle = 29$. After one step, it has four equally likely positions:

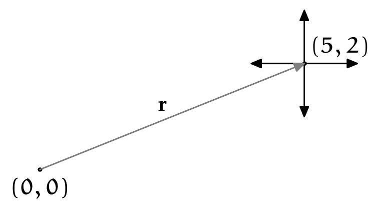

Rather than compute the new expected squared distance using all four positions, be lazy and just look at the two horizontal motions. The two possibilities are $(6, 2)$ and $(4, 2)$. The expected squared distance is

$$\langle x^2 \rangle = \frac{1}{2}(40 + 20) = 30,$$

which is one more than the previous value of $\langle x^2 \rangle$. Since nothing is special about horizontal motion compared to vertical motion – symmetry! – the same result holds for vertical motion. So, averaging over the four possible locations produces an expected squared distance of 30.

For two dimensions, the pattern is:

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

No step in the analysis depended on being in only two dimensions. In fancy words, the derivation and the result are invariant to change of dimensionality. In plain English, this result also works in three dimensions.

In a standard walk in a straight line, $\langle x \rangle \propto$ time. Note the single power of x. The interesting quantity in a regular walk is not x itself, since it can grow without limit and is not invariant, but the ratio x/t, which is invariant to changes in t. This invariant is known as the speed.

In a random walk, where $\langle x^2 \rangle \propto t$, the interesting quantity is $\langle x^2 \rangle / t$. The expected squared position is not invariant to changes in t. However, the ratio $\langle x^2 \rangle / t$ is invariant. This invariant is, except for a dimensionless constant, the *diffusion constant* and is often denoted D. It has dimensions of L^2T^{-1}.

This qualitative difference between a random and a regular walk makes intuitive sense. A random walker, for example a gas molecule or a very

So essenially on average for a 2 dimensional random walk, the most likely probability is a ring that is slowly moving outwards?

Actually the most likely position is the origin! The distribution is roughly Gaussian (a normal distribution) with variance <x^2>. So the distribution has a peak at the origin, and it spreads farther and farther with time.

I like this simplification- I think it will help me to understand the example in a more simple way while still being able to apply the rules to a more complex example

so we limited it to horizontal but later we don't really go back and zoom out to include vertical –> shouldn't this change the answer?

i'm confused as to where these numbers are coming from.

From the vector coordinates: 40=6^2 + 2^2 20=4^2 + 2^2

Thanks! This is helpful.

I guessed this but I think it would be helpful to include somewhere before here how <x^2> will be calculated in 2D.

This is pretty cool that it works out like this again. The notation confused me a bit but this helps clear things up.

I like symmetry–makes things easy.

And it seems to be much easier to apply when you're talking about probabilities and can simply take the average.

I'm not sure why, but when you mention symmetry this example seems more effective than the previous.

now

I like this paragraph! even if the 1st sentence is a little awkward. it's great.

This paragraph reads weird when you understand all three sentences, but it is a thorough explanation.

The result that $\langle x^2\rangle$ is proportional to time applies not only to the one-dimensional random walk. Here's an example in two dimensions. Suppose that the particle's position is $(5,2)$, so $\langle x^2\rangle = 29$. After one step, it has four equally likely positions:

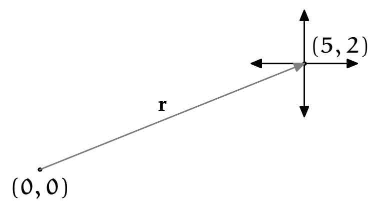

Rather than compute the new expected squared distance using all four positions, be lazy and just look at the two horizontal motions. The two possibilities are $(6,2)$ and $(4,2)$. The expected squared distance is

$$\langle x^2\rangle = \frac{1}{2}(40 + 20) = 30,$$

which is one more than the previous value of $\langle x^2\rangle$. Since nothing is special about horizontal motion compared to vertical motion – symmetry! – the same result holds for vertical motion. So, averaging over the four possible locations produces an expected squared distance of 30.

For two dimensions, the pattern is:

$$\langle x^2_{n+1}\rangle = \langle x^2_n\rangle + 1.$$

No step in the analysis depended on being in only two dimensions. In fancy words, the derivation and the result are invariant to change of dimensionality. In plain English, this result also works in three dimensions.

In a standard walk in a straight line, $\langle x\rangle \propto$ time. Note the single power of x. The interesting quantity in a regular walk is not x itself, since it can grow without limit and is not invariant, but the ratio x/t, which is invariant to changes in t. This invariant is known as the speed.

In a random walk, where $\langle x^2\rangle \propto t$, the interesting quantity is $\langle x^2\rangle/t$. The expected squared position is not invariant to changes in t. However, the ratio $\langle x^2\rangle/t$ is invariant. This invariant is, except for a dimensionless constant, the *diffusion constant* and is often denoted D. It has dimensions of $\mathrm{L}^2\mathrm{T}^{-1}$.

This qualitative difference between a random and a regular walk makes intuitive sense. A random walker, for example a gas molecule or a very

This is a great sentence. Very clear

I think saying this 2 different ways is unnecessary

Interesting to see speed as an invariant.

I am not sure that I understand why, unless we made the assumption of constant motion previously.

Agreed, I've always seen it as an easy and common variable to change.

How is this invariant?

I had the same question, a proof of this would be nice.

it's invariant because we're looking at distance/speed = velocity, and one of our assumptions was that the molecule moves at a constant speed.

x/t is invariant but x is variant because if we look at x, it depends on where we just were. but the rate theoretically is always the same. like if i walked 3 steps, i could walk forward forward forward or forward backward forward, and my x would be different. but the pace at which i walk would still be the same

i didn't catch this distinction here.

This section is kind of confusing. I don't see how <x^2> is variant with different t but the ratio is invariant.

The <x^2> is variant because it always changes where you were previously before the step. However, the rate/speed at which you move is always the same. I like to think that I'm walking one-step/second.

The <x^2> is variant because it always changes where you were previously before the step. However, the rate/speed at which you move is always the same. I like to think that I'm walking one-step/second.

If x = 2*t, then x is not invariant to changes in t (t goes up, x goes up). However, x/t is always 2 - an invariant.

It feels like this paragraph is really dense...

The result that $\langle x^2 \rangle$ is proportional to time applies not only to the one-dimensional random walk. Here's an example in two dimensions. Suppose that the particle's position is $(5, 2)$, so $\langle x^2 \rangle = 29$. After one step, it has four equally likely positions:

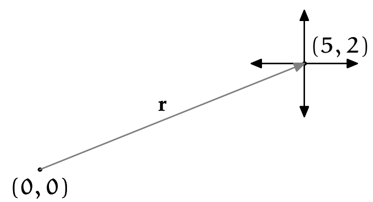

Rather than compute the new expected squared distance using all four positions, be lazy and just look at the two horizontal motions. The two possibilities are $(6, 2)$ and $(4, 2)$. The expected squared distance is

$$\langle x^2 \rangle = \frac{1}{2}(40 + 20) = 30,$$

which is one more than the previous value of $\langle x^2 \rangle$. Since nothing is special about horizontal motion compared to vertical motion – symmetry! – the same result holds for vertical motion. So, averaging over the four possible locations produces an expected squared distance of 30.

For two dimensions, the pattern is:

$$\langle x_{n+1}^2 \rangle = \langle x_n^2 \rangle + 1.$$

No step in the analysis depended on being in only two dimensions. In fancy words, the derivation and the result are invariant to change of dimensionality. In plain English, this result also works in three dimensions.

In a standard walk in a straight line, $\langle x \rangle \propto$ time. Note the single power of x. The interesting quantity in a regular walk is not x itself, since it can grow without limit and is not invariant, but the ratio x/t, which is invariant to changes in t. This invariant is known as the speed.

In a random walk, where $\langle x^2 \rangle \propto t$, the interesting quantity is $\langle x^2 \rangle / t$. The expected squared position is not invariant to changes in t. However, the ratio $\langle x^2 \rangle / t$ is invariant. This invariant is, except for a dimensionless constant, the *diffusion constant* and is often denoted D. It has dimensions of $\mathrm{L}^2\mathrm{T}^{-1}$.

This qualitative difference between a random and a regular walk makes intuitive sense. A random walker, for example a gas molecule or a very

Oh wow, even though I should have known we were going here eventually, this caught me completely off-guard by being perfectly intuitive. x/t –> x^2/t. Diffusion, makes sense.

You might note, more explicitly, that this doesn't depend on whether we are in the 1D, 2D, or higher-dimensional case. (I had to go back and double check the math to convince myself.)

Agreed - you might want to make the conclusion more explicit.

He says above that the result is invariant to change of dimensionality.

I feel like this point was made pretty clear shortly before this.

This is pretty cool, I've never seen any of this before and it's definitely not something intuitive

this has units of length * speeed. but i thought the invariant WAS speed. could you explain please?

The invariant is speed in the case of the normal walk, but for the random walk the diffusionconstant is the invariant.

We've just changed cases and selected the next appropriate invariant. From x/t to x^2/t

I've definitely seen easier ways to explain the difference between regular and random walks; however, i think your explanation is more relevant to the course.

I like seeing what would classify a random walker, it helps to get a feel for what we are really looking at here.

I have a feeling it's anything that could be classified in terms of the motion described in this section.

I like the two completely different examples reinforcing the same point.

drunk person, moves back and forth, sometimes making progress in one direction, and other times undoing that progress. So, in order to travel the same distance, a random walker should require longer than a regular walker requires. The relation $\langle x^2 \rangle / t \sim D$ confirms and sharpens this intuition. The time for a random walker to travel a distance l is $t \sim l^2/D$, which grows quadratically rather than linearly with distance.

8.3.2 Diffusion equation

The preceding conclusion about random walks is sufficient to derive the diffusion equation, which describes how charge (electrons) move in a wire, how heat conducts through solid objects, and how gas molecules travel. Imagine then a gas of particles with each particle doing a random walk in one dimension. What is the equation that describes how the concentration, or number density, varies with time?

Divide the one-dimensional world into slices of width Δx, where Δx is the mean free path. Then look at the slices at $x - \Delta x$, x, and $x + \Delta x$. In every time step, one-half the molecules in each slice move left, and one-half move right. So the number of molecules in the x slice changes from $N(x)$ to

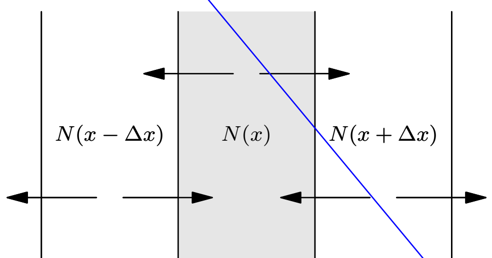

$$\frac{1}{2}(N(x - \Delta x) + N(x + \Delta x)).$$

The change in N is

$$\Delta N = \frac{1}{2}(N(x - \Delta x) + N(x + \Delta x)) - N(x)$$
$$= \frac{1}{2}(N(x - \Delta x) - 2N(x) + N(x + \Delta x)).$$

This last relation can be rewritten as

$$\Delta N \sim (N(x + \Delta x) - N(x)) - (N(x) - N(x + \Delta x))\,.$$

In terms of derivatives, it is

$$\Delta N \sim (\Delta x)^2 \frac{\partial^2 N}{\partial x^2}.$$

Comments on page 4

nice example :)

I was wondering where this would come in when it was missing from that intro paragraph...

haha

do you mean distance or displacement?

Except in the (unlikely) case that the drunk and regular walker make the same moves.

This is why he uses the word 'should'....but yes this is possible as well.

is this valid only when l is very small

I don't think so. It still makes sense over longer periods of time.

This a pretty interesting statement. However, it intuively doesn't make sense because on average, the drunk person would have still traveled zero distance.

Yes, but the expected value of the square is different, since in the case where the mean is zero, that is the variance, which would be nonzero in this case.

well also the case that he moved exactly nowhere at any given point seems intuitively unlikely, although expected value of his overall movement may still be zero.

This of it this way. On average, he would make zero progress, but randomly he'll one way make his way home. Now if he's so drunk he doesn't recognize his own house, he might go all the way back and cancel out that distance, but we assume that once he somehow reaches his house, he goes inside.

I really like these explanations using intuition.

I think these intuitive explanations solidify my understanding.

This section was very interesting - it made a lot of other things I've seen before about moments make a lot more sense.

drunk person, moves back and forth, sometimes making progress in one direction, and other times undoing that progress. So, in order to travel the same distance, a random walker should require longer than a regular walker requires. The relation $\langle x^2\rangle/t \sim D$ confirms and sharpens this intuition. The time for a random walker to travel a distance l is $t \sim l^2/D$, which grows quadratically rather than linearly with distance.

8.3.2 Diffusion equation

The preceding conclusion about random walks is sufficient to derive the diffusion equation, which describes how charge (electrons) move in a wire, how heat conducts through solid objects, and how gas molecules travel. Imagine then a gas of particles with each particle doing a random walk in one dimension. What is the equation that describes how the concentration, or number density, varies with time?

Divide the one-dimensional world into slices of width Δx, where Δx is the mean free path. Then look at the slices at $x - \Delta x$, x, and $x + \Delta x$. In every time step, one-half the molecules in each slice move left, and one-half move right. So the number of molecules in the x slice changes from $N(x)$ to

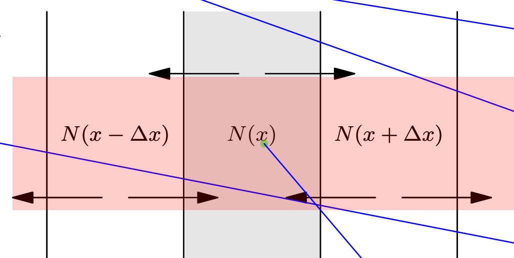

$$\frac{1}{2}(N(x-\Delta x)+N(x+\Delta x)).$$

The change in N is

$$\Delta N = \frac{1}{2}(N(x-\Delta x)+N(x+\Delta x)) - N(x)$$
$$= \frac{1}{2}(N(x-\Delta x)-2N(x)+N(x+\Delta x)).$$

This last relation can be rewritten as

$$\Delta N \sim (N(x+\Delta x)-N(x)) - (N(x)-N(x+\Delta x))\,.$$

In terms of derivatives, it is

$$\Delta N \sim (\Delta x)^2 \frac{\partial^2 N}{\partial x^2}.$$

Would this change if we were observing in three dimensions. For example if we were discussing an airplane's movements or a spacecraft.

It's the same conclusion in any number of dimensions: the variance grows linearly with time. I'll give another argument in class to show it by using dot products. The argument is more general simpler than the one in the text, but somehow less convincing (at least to me) because it is algebraic and not numerical. I think I'll put both arguments in the text, so that people of different persuasions can be persuaded by whatever resonates.

Again, it's nice to see examples of what we are going to be talking about before it is introduced.

Instead of introducing this section like this, I think it might be helpful to make the transition by explaining that random walks are simple mathematical models of diffusion.

I like how we eased into this section with simple explanations before diving into more complex equations.

I really like how this complex example relates back to what we just did with random walks- I'm glad to see the two separated into 2 examples rather than 1 example that crams both in.

what does this mean?

= how far on average the particle travels before it hits something or is forced to change its trajectory

Thanks for the clarification, I actually initially interpreted this phrase differently but that makes more sense.

ah...this diagram reminds me of 8.02 using Gauss' Law.

Haha, it does for me too. This I think is a fantastic diagram!

this diagram is a little confusing? what are the three layers?

This is a much simpler but nicer explanation of what was covered in another class (which did not, unfortunately, let us calculate using estimation!)

drunk person, moves back and forth, sometimes making progress in one direction, and other times undoing that progress. So, in order to travel the same distance, a random walker should require longer than a regular walker requires. The relation $\langle x^2 \rangle / t \sim D$ confirms and sharpens this intuition. The time for a random walker to travel a distance l is $t \sim l^2/D$, which grows quadratically rather than linearly with distance.

8.3.2 Diffusion equation

The preceding conclusion about random walks is sufficient to derive the diffusion equation, which describes how charge (electrons) move in a wire, how heat conducts through solid objects, and how gas molecules travel. Imagine then a gas of particles with each particle doing a random walk in one dimension. What is the equation that describes how the concentration, or number density, varies with time?

Divide the one-dimensional world into slices of width Δx, where Δx is the mean free path. Then look at the slices at $x - \Delta x$, x, and $x + \Delta x$. In every time step, one-half the molecules in each slice move left, and one-half move right. So the number of molecules in the x slice changes from $N(x)$ to

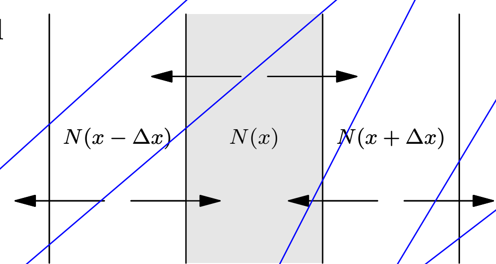

$$\frac{1}{2}(N(x - \Delta x) + N(x + \Delta x)).$$

The change in N is

$$\Delta N = \frac{1}{2}(N(x - \Delta x) + N(x + \Delta x)) - N(x)$$
$$= \frac{1}{2}(N(x - \Delta x) - 2N(x) + N(x + \Delta x)).$$

This last relation can be rewritten as

$$\Delta N \sim (N(x + \Delta x) - N(x)) - (N(x) - N(x + \Delta x)).$$

In terms of derivatives, it is

$$\Delta N \sim (\Delta x)^2 \frac{\partial^2 N}{\partial x^2}.$$

I feel as though this concept in particular is best explained in class.

Is this always true in every dimension?

I think this is true in every direction because as long as we assume similar conditions, diffusion can happen in any direction.

Tracing this particular jump in logic was a little hard.

Is this supposed to be -N(x- delta x)?

Otherwise, the expression reduces to 0...

Well, no, it doesn't, it reduces to 2N(x+dx)-2N(x), but I'm pretty sure it is supposed to be +N(x-dx).

Whoops, thanks for catching that mistake.

Thank you for explaining how you got this. I would have had to take it at face value if you didn't.

Where does he explain this?

If I saw this equation before you broke it down, I would be terrified. Thanks for making it seem simple

I agree; the explanations prior to introducing the equation make it a lot clearer.

I think most of us have seen the technical expression for a derivative as a limit, but probably fewer have seen the connection of a second derivative to a limit. This is expression was not an obvious result of the previous line, and it took me some time to convince myself that it was reasonable...

You're right, I should explain that (and will do so in class).

The slices are separated by a distance such that most of the molecules travel from one piece to the neighboring piece in the time step τ. If τ is the time between collisions – the mean free time – then the distance is the mean free path λ. Thus

$$\frac{\Delta N}{\tau} \sim \frac{\lambda^2}{\tau}\frac{\partial^2 N}{\partial x^2},$$

or

$$\dot{N} \sim D\frac{\partial^2 N}{\partial x^2}$$

where $D \sim \lambda^2/\tau$ is a diffusion constant.

This partial-differential equation has interesting properties. The second spatial derivative means that a linear spatial concentration gradient remains unchanged. Its second derivative is zero so its time derivative must be zero. Diffusion fights curvature – roughly speaking, the second derivative – and does not try to change the gradient directly.

8.3.3 Keeping warm

One consequence of the diffusion equation is an analysis of how to keep warm on a cold day. To quantify keeping warm, or feeling cold, we need to calculate the heat flux: the energy flowing per unit area per unit time. Start with the definition of flux. Flux (of anything) is defined as

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{area} \times \text{time}}.$$

The flux depends on the density of stuff and on how fast the stuff travels:

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{volume}} \times \text{speed}.$$

For heat flux, the stuff is thermal energy. The specific heat c_p is the thermal energy per mass per temperature, c_pT is the thermal energy per mass, and $\rho c_p T$ is therefore the thermal energy per volume. The speed is the 'speed' of diffusion. To diffuse a distance l takes time $t \sim l^2/D$, making the speed l/t or D/l. The l in the denominator indicates that, as expected, diffusion is slow over long distances. For heat diffusion, the diffusion constant is denoted κ and called the thermal diffusivity. So the speed is l/κ.

how does this correlate to the mean free path?

This may be intuitive to all of the course 2 majors out there but as a course 6 major, I find myself getting very confused about all this terminology I haven't seen before.

It helps to look at it a bit more mathematically and work yourself through the problem moving units around as they should be.

Thanks! Thinking about the units did help clarify where all these quantities come from.

that "dot" is kind of hard to see. I thought it was a piece of dust on my screen at first. Maybe use dN/dt?

Good point.

David Hogg, a friend who teaches at NYU, said he doesn't even allow his students to use the dot notation because it obscures the dimensions, whereas dN/dt makes the dimensions clear. So he recommended that I take out all uses of the dot notation from _Street-Fighting Mathematics_. Which I did, and I'll do the same for this book.

I feel like I've seen this in a microelectronics class. Is this the same principle used in calculating thermal equilibrium states in pn junctions?

Oh, that makes sense now.

Oh man, this looks suspiciously like the stuff we did above with the second moment. I like it.

what are some values of D to give us a feel for what it is

For gas molecules diffusing around (e.g. air molecules), D is about 10^-5 m^2/s – which is the kinematic viscosity of air.

Is it really a constant if tau and gamma are variable for each problem? I think of a constant as something like Avogadro's number which never changes..

This paragraph goes through a lot of information very quickly, it might be worth it to explain it more.

The slices are separated by a distance such that most of the molecules travel from one piece to the neighboring piece in the time step τ. If τ is the time between collisions – the mean free time – then the distance is the mean free path λ. Thus

$$\frac{\Delta N}{\tau} \sim \frac{\lambda^2}{\tau}\frac{\partial^2 N}{\partial x^2},$$

or

$$\dot{N} \sim D\frac{\partial^2 N}{\partial x^2}$$

where $D \sim \lambda^2/\tau$ is a diffusion constant.

This partial-differential equation has interesting properties. The second spatial derivative means that a linear spatial concentration gradient remains unchanged. Its second derivative is zero so its time derivative must be zero. Diffusion fights curvature – roughly speaking, the second derivative – and does not try to change the gradient directly.

8.3.3 Keeping warm

One consequence of the diffusion equation is an analysis of how to keep warm on a cold day. To quantify keeping warm, or feeling cold, we need to calculate the heat flux: the energy flowing per unit area per unit time. Start with the definition of flux. Flux (of anything) is defined as

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{area} \times \text{time}}.$$

The flux depends on the density of stuff and on how fast the stuff travels:

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{volume}} \times \text{speed}.$$

For heat flux, the stuff is thermal energy. The specific heat c_p is the thermal energy per mass per temperature, c_pT is the thermal energy per mass, and ρc_pT is therefore the thermal energy per volume. The speed is the 'speed' of diffusion. To diffuse a distance l takes time $t \sim l^2/D$, making the speed l/t or D/l. The l in the denominator indicates that, as expected, diffusion is slow over long distances. For heat diffusion, the diffusion constant is denoted κ and called the thermal diffusivity. So the speed is l/κ.

I find that this is the easiest way for me to remember how (thermal, in my context) diffusivity acts. It might be worth elaborating on.

i.e. If a long thin rod has a heat source at one end set to 50 degrees, and a heat sink at the other end set to 0 degrees, diffusion will act to make the temperature profile a straight line (from 50 at one end to 0 at the other) rather than some quadratic curve, etc.

thank you for the explanation! i was just about to ask what "diffusion fights curvature" means.

Haha well explained! Again, had someone in my course 20 class explained it like that it would have made life a lot easier...This explanation might actually be worthwhile to add into the book

yes, this was great! Even to a non-engineering or science major, I was able to clearly understand this property with the above example.

I had almost exactly that explanation but without a diagram. Instead of making a diagram (it was late at night) I decided to take out the example – clearly the wrong choice!

it keeps my attention when you relate it to a real world example the whole way through instead of briefly in introductory sentences.

so this section actually doesn't have to do with random walks. it just stems from the diffusion equation. you should make that more clear in the intro.

He mentions that after he concludes on the diffusion constant.

The diffusion equation is the macroscopic view of heat flow. Under the hood, it is a random walk that produces the diffusion equation.

is there a way to quantify how some people get colder easier than others?

I now understand the sentence in the very first paragraph about keeping warm, I didn't realize then that you were talking about diffusion. I now understand why "keeping warm" is part of random walks.

the day i learned about fins in 2.005 was the day i understood why my family (all impossibly tall and lanky) is perennially cold.

The slices are separated by a distance such that most of the molecules travel from one piece to the neighboring piece in the time step τ. If τ is the time between collisions – the mean free time – then the distance is the mean free path λ. Thus

$$\frac{\Delta N}{\tau} \sim \frac{\lambda^2}{\tau}\frac{\partial^2 N}{\partial x^2},$$

or

$$\dot{N} \sim D\frac{\partial^2 N}{\partial x^2}$$

where $D \sim \lambda^2/\tau$ is a diffusion constant.

This partial-differential equation has interesting properties. The second spatial derivative means that a linear spatial concentration gradient remains unchanged. Its second derivative is zero so its time derivative must be zero. Diffusion fights curvature – roughly speaking, the second derivative – and does not try to change the gradient directly.

8.3.3 Keeping warm

One consequence of the diffusion equation is an analysis of how to keep warm on a cold day. To quantify keeping warm, or feeling cold, we need to calculate the heat flux: the energy flowing per unit area per unit time. Start with the definition of flux. Flux (of anything) is defined as

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{area} \times \text{time}}.$$

The flux depends on the density of stuff and on how fast the stuff travels:

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{volume}} \times \text{speed}.$$

For heat flux, the stuff is thermal energy. The specific heat c_p is the thermal energy per mass per temperature, c_pT is the thermal energy per mass, and ρc_pT is therefore the thermal energy per volume. The speed is the 'speed' of diffusion. To diffuse a distance l takes time $t \sim l^2/D$, making the speed l/t or D/l. The l in the denominator indicates that, as expected, diffusion is slow over long distances. For heat diffusion, the diffusion constant is denoted κ and called the thermal diffusivity. So the speed is l/κ.

I would have loved to use approaches this elegant and stream-lined in 2.005

I don't think this needs to be simplified to the point of using the word "stuff" - I actually find that more confusing

i think he means to generalize the def of flux, so it's the flux of ANYTHING. therefore "stuff"

I agree with the first comment. I find the word "stuff" confusing. We've covered much more confusing information in this class, so I think making it a bit more complicated than "stuff" wouldn't be a problem.

What are the dimensions of stuff? M?

"Stuff" refers to idea that many things can represent that variable. For example, [stuff]=M if we were looking at mass-flux but it could also [stuff]=energy in energy-flux and [stuff]=L^3 for volumetric-flux.

I think this is helpful because it is a general formula for flux–stuff can mean anything. the general idea is how much of something flows through a given area per time

I was trying to make it more general. In solar flux, for example, the stuff is energy; and flux of stuff is energy/area/time or power/area. There's also momentum flux, where the stuff is momentum; and number (or partcile) flux, where the stuff is particles. All of those definitions of flux have the same structure (an abstraction!), and I wanted to make that explicit.

Perhaps in the text I should give one or two other examples of using the abstraction.

I like this very simplified explanation.

Yeah it's simple, but the best part is that it actually does apply to a variety of things and makes sense at the same time.

The slices are separated by a distance such that most of the molecules travel from one piece to the neighboring piece in the time step τ. If τ is the time between collisions – the mean free time – then the distance is the mean free path λ. Thus

$$\frac{\Delta N}{\tau} \sim \frac{\lambda^2}{\tau}\frac{\partial^2 N}{\partial x^2},$$

or

$$\dot{N} \sim D\frac{\partial^2 N}{\partial x^2}$$

where $D \sim \lambda^2/\tau$ is a diffusion constant.

This partial-differential equation has interesting properties. The second spatial derivative means that a linear spatial concentration gradient remains unchanged. Its second derivative is zero so its time derivative must be zero. Diffusion fights curvature – roughly speaking, the second derivative – and does not try to change the gradient directly.

8.3.3 Keeping warm

One consequence of the diffusion equation is an analysis of how to keep warm on a cold day. To quantify keeping warm, or feeling cold, we need to calculate the heat flux: the energy flowing per unit area per unit time. Start with the definition of flux. Flux (of anything) is defined as

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{area} \times \text{time}}.$$

The flux depends on the density of stuff and on how fast the stuff travels:

$$\text{flux of stuff} = \frac{\text{stuff}}{\text{volume}} \times \text{speed}.$$

For heat flux, the stuff is thermal energy. The specific heat c_p is the thermal energy per mass per temperature, c_pT is the thermal energy per mass, and $\rho c_p T$ is therefore the thermal energy per volume. The speed is the 'speed' of diffusion. To diffuse a distance l takes time $t \sim l^2/D$, making the speed l/t or D/l. The l in the denominator indicates that, as expected, diffusion is slow over long distances. For heat diffusion, the diffusion constant is denoted κ and called the thermal diffusivity. So the speed is l/κ.

i find this easier to understand as [(stuff per area) per time].

For me it's just the opposite. I find (stuff per time, per area) slightly more intuitive. I guess this is a good compromise.

I like both of these definitions because they both describe this equation better than the reading.

This makes a lot of sense to me. I think it is because flux is such an important concept in classes like 8.02.

I like this ... I think it's very well written.

volume of stuff? or volume stuff is flowing into?

Per unit volume... stuff is flowing into.

This note is going to a direction that I had not imagined; I didn't think we would start with random walk and come here.

I'm not sure I follow how we got this.

That's a neat, simple, and convincing justification.

It makes sense that over a long period of time, little diffusion occurs. Taking for instance heat. When you apply the same amount of heat to an object, the change in temperature at the beginning is more than the change in temperature at a later time in the same time interval

kappa/l ?

yes. he definitely has it inverted.

You're right. Thanks – I should check the dimensions of everything that I write, not only teach other people to do the same.

Combine the thermal energy per volume with the diffusion speed:

$$\text{thermal flux} = \rho c_p T \times \frac{\kappa}{l}.$$

The product $\rho c_p \kappa$ occurs so frequently that it is given a name: the thermal conductivity K. The ratio T/l is a lumped version of the temperature gradient $\Delta T/\Delta x$. With those substitutions, the thermal flux is

$$F = K\frac{\Delta T}{\Delta x}.$$

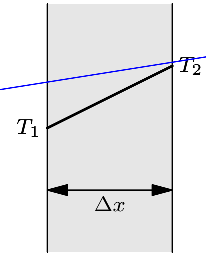

With one side held at T_1 and the other at T_2, the temperature gradient is $(T_2 - T_1)/\Delta x$.

To estimate how much heat one loses on a cold day, we need to estimate $K = \rho c_p \kappa$. To do so, put all the pieces together:

$$\rho \sim 1\,\text{kg}\,\text{m}^{-3},$$
$$c_p \sim 10^3\,\text{J}\,\text{kg}^{-1}\,\text{K}^{-1},$$
$$\kappa \sim 1.5\cdot 10^{-5}\,\text{m}^2\,\text{s}^{-1},$$

where we are guessing that $\kappa = \nu$ (because both κ and ν are diffusion constants). Then

$$K = \rho c_p \kappa \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1}.$$

Using this value we can estimate the heat loss on a cold day. Let's say that your skin is at $T_2 = 30\,°\text{C}$ and the air outside is $T_1 = 0\,°\text{C}$, making $\Delta T = 30\,\text{K}$. A thin T-shirt may have thickness 2 mm, so

$$F = K\frac{\Delta T}{\Delta x} \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1} \times \frac{30\,\text{K}}{2\cdot 10^{-3}\,\text{m}} \sim 300\,\text{W}\,\text{m}^{-2}.$$

Damn, we want a power rather than a power per area. Ah, flux is power per area, so just multiply by a person's surface area: roughly 2 m tall and 0.5 m wide, with a front and a back. So the surface area is about $2\,\text{m}^2$. Thus, the power lost is

$$P \sim FA = 300\,\text{W}\,\text{m}^{-2} \times 2\,\text{m}^2 = 600\,\text{W}.$$

No wonder a winter day wearing only thin pants and shirt feels so cold: 600 W is large compared to human power levels. Sitting around, a person produces 100 W of heat (the basal metabolic rate). When 600 W escapes,

It would be easier to have this equation on the last page, so we don't flip from page to page to see how this equation was derived.

I still find it amazing in science and math in general that so many formulas can condense down into one simple equation, with lots of lumped "black boxes" in them (even if in this case we know what the constant is hiding)

I like the explanation. I found it easy to follow the path taken.

Agreed. It was much easier to comprehend than the previous sections.

I found the last section pretty comprehensible, but I agree that this one was very elegant and well written.

That's a benefit of reading memos. When I first taught a version of this class (in IAP 2006), the class did reading memos on paper, which they turned in the day we had the lecture on the topic. So, I couldn't use the reading memos to fix that particular chapter in time to give students (since they had already read it).

But I did use the previous memos to figure out how students were thinking and what explanations would help. With this particular subsection, I remember using what I learnt from the previous memos and then spending 8 hours straight rewriting the explanation of flux and diffusion until I was happy with it.

how can we make this assumption? what is k again?

K is the thermal diffusivity, nu is the kinematic viscosity (momentum diffusivity constant).

I think that the model you are using for heat loss needs to be more clearly explained before diving into all of these quantities. For instance, it appears you are calculating the K for air, but you don't say so (or why) before jumping right in. I think that'd be useful and would make this analysis clearer.

I agree, some of us have not taken that much thermo-stuff, and some background would definitly help.

How does wind affect this calculation? Do you just insert the wind-chill temperature in for T1?

Combine the thermal energy per volume with the diffusion speed:

$$\text{thermal flux} = \rho c_p T \times \frac{\kappa}{l}.$$

The product $\rho c_p \kappa$ occurs so frequently that it is given a name: the thermal conductivity K. The ratio T/l is a lumped version of the temperature gradient $\Delta T/\Delta x$. With those substitutions, the thermal flux is

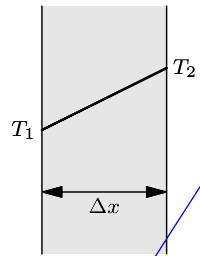

$$F = K\frac{\Delta T}{\Delta x}.$$

With one side held at T_1 and the other at T_2, the temperature gradient is $(T_2 - T_1)/\Delta x$.

To estimate how much heat one loses on a cold day, we need to estimate $K = \rho c_p \kappa$. To do so, put all the pieces together:

$$\rho \sim 1\,\text{kg}\,\text{m}^{-3},$$
$$c_p \sim 10^3\,\text{J}\,\text{kg}^{-1}\,\text{K}^{-1},$$
$$\kappa \sim 1.5\cdot 10^{-5}\,\text{m}^2\,\text{s}^{-1},$$

where we are guessing that $\kappa = \nu$ (because both κ and ν are diffusion constants). Then

$$K = \rho c_p \kappa \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1}.$$

Using this value we can estimate the heat loss on a cold day. Let's say that your skin is at $T_2 = 30\,°\text{C}$ and the air outside is $T_1 = 0\,°\text{C}$, making $\Delta T = 30\,\text{K}$. A thin T-shirt may have thickness 2 mm, so

$$F = K\frac{\Delta T}{\Delta x} \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1} \times \frac{30\,\text{K}}{2\cdot 10^{-3}\,\text{m}} \sim 300\,\text{W}\,\text{m}^{-2}.$$

Damn, we want a power rather than a power per area. Ah, flux is power per area, so just multiply by a person's surface area: roughly 2 m tall and 0.5 m wide, with a front and a back. So the surface area is about $2\,\text{m}^2$. Thus, the power lost is

$$P \sim FA = 300\,\text{W}\,\text{m}^{-2} \times 2\,\text{m}^2 = 600\,\text{W}.$$

No wonder a winter day wearing only thin pants and shirt feels so cold: 600 W is large compared to human power levels. Sitting around, a person produces 100 W of heat (the basal metabolic rate). When 600 W escapes,

How does wind affect this calculation? Do you just insert the wind-chill temperature in for T1?

I'd assume so, since when there's wind chill it "feels like" some lower temperature.

Wind chill is a very interesting application of all the previous ideas. The wind chill affects your exposed skin (the wind doesn't matter to the areas protected by a jacket). On the exposed areas, a fast wind means a thin boundary layer (shorter time for momentum to diffuse), so the temperature gradient DT/Dx from skin to air is large – hence there is more heat flowing from your skin, just as it would on a windless but colder day.

Good question, I'd be curious to know the answer of this.

is this supposed to be the thermal conductivity of air? the standard values for that are 10-100 W/mK.

Typical solids are about 1 W/(m*K) and metals (due to electrons being good movers of heat) are around 100 W/(m*K) but air is a lot lower than all of the above. Wikipedia gives 0.025 W/(m*K).

(To the original poster) Could you be thinking about convection over a surface (often h=10 to 100 W/m^2K) rather than conduction?

It's pretty cool deriving and using the equations we derived months ago in 2.005 in a completely different way!

Seriously, I wonder if there are equations that physics hasn't been able to simplify but these types of methods might be able bring us to some kind of close approximation.

Combine the thermal energy per volume with the diffusion speed:

$$\text{thermal flux} = \rho c_p T \times \frac{\kappa}{l}.$$

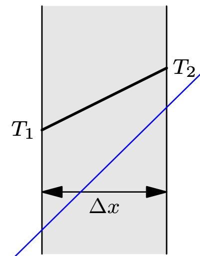

The product $\rho c_p \kappa$ occurs so frequently that it is given a name: the thermal conductivity K. The ratio T/l is a lumped version of the temperature gradient $\Delta T/\Delta x$. With those substitutions, the thermal flux is

$$F = K\frac{\Delta T}{\Delta x}.$$

With one side held at T_1 and the other at T_2, the temperature gradient is $(T_2 - T_1)/\Delta x$.

To estimate how much heat one loses on a cold day, we need to estimate $K = \rho c_p \kappa$. To do so, put all the pieces together:

$$\rho \sim 1\,\text{kg}\,\text{m}^{-3},$$
$$c_p \sim 10^3\,\text{J}\,\text{kg}^{-1}\,\text{K}^{-1},$$
$$\kappa \sim 1.5 \cdot 10^{-5}\,\text{m}^2\,\text{s}^{-1},$$

where we are guessing that $\kappa = \nu$ (because both κ and ν are diffusion constants). Then

$$K = \rho c_p \kappa \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1}.$$

Using this value we can estimate the heat loss on a cold day. Let's say that your skin is at $T_2 = 30\,°\text{C}$ and the air outside is $T_1 = 0\,°\text{C}$, making $\Delta T = 30\,\text{K}$. A thin T-shirt may have thickness 2 mm, so

$$F = K\frac{\Delta T}{\Delta x} \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1} \times \frac{30\,\text{K}}{2 \cdot 10^{-3}\,\text{m}} \sim 300\,\text{W}\,\text{m}^{-2}.$$

Damn, we want a power rather than a power per area. Ah, flux is power per area, so just multiply by a person's surface area: roughly 2 m tall and 0.5 m wide, with a front and a back. So the surface area is about $2\,\text{m}^2$. Thus, the power lost is

$$P \sim FA = 300\,\text{W}\,\text{m}^{-2} \times 2\,\text{m}^2 = 600\,\text{W}.$$

No wonder a winter day wearing only thin pants and shirt feels so cold: 600 W is large compared to human power levels. Sitting around, a person produces 100 W of heat (the basal metabolic rate). When 600 W escapes,

I'm not sure this is appropriate for a textbook....

It is Sanjoy's textbook...

yeah, Sanjoy knows we are all adults anyway

I think this situation doesn't warranty that extreme of a response. It seems like a minor annoyance as opposed to something you would curse about. (irregardless of whether it is appropriate or not)

lol I'm surprised anybody actually commented on this

I think it's totally appropriate – it makes it seem like a person is speaking to you instead of some black and white text on a page is proclaiming something.

HAHAAHAHAHAHAHAHA. Love it

I mean I think it's fine... This whole textbook is pretty conversational.

I feel like this math was really well done. thanks

Hmm, I think at 600W power loss, your internal temperature drop would be significant, no? (i.e. your skin temperature would be significantly low). This would lower our result, maybe a significant amount (perhaps expand the example to include skin and flesh as thermal insulators?)

Right, if you didn't do anything about it. Hence your body starts shivering, and it's almost an unconscious action – to burn enough fuel to keep your core temperature constant. Your body is willing to let the extremities get too cold, if it has to (hence frostbite) in order to keep the core temperature from dropping too much. Enzymes have very tightly optimized shapes and even small changes in temperature can significantly lower reaction rates.

wow...

I've always heard that your head and feet lose the most heat...is that true and how does that fact play into this equation?

Yeah that's true. I wonder if that plays a role here, or if it all averages out to something close to our estimate

Combine the thermal energy per volume with the diffusion speed:

$$\text{thermal flux} = \rho c_\text{p} T \times \frac{\kappa}{l}.$$

The product $\rho c_\text{p} \kappa$ occurs so frequently that it is given a name: the thermal conductivity K. The ratio T/l is a lumped version of the temperature gradient $\Delta T/\Delta x$. With those substitutions, the thermal flux is

$$F = K\frac{\Delta T}{\Delta x}.$$

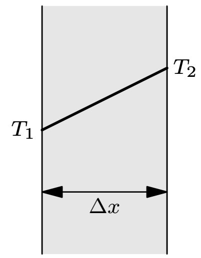

With one side held at T_1 and the other at T_2, the temperature gradient is $(T_2 - T_1)/\Delta x$.

To estimate how much heat one loses on a cold day, we need to estimate $K = \rho c_\text{p} \kappa$. To do so, put all the pieces together:

$$\rho \sim 1\,\text{kg}\,\text{m}^{-3},$$
$$c_\text{p} \sim 10^3\,\text{J}\,\text{kg}^{-1}\,\text{K}^{-1},$$
$$\kappa \sim 1.5 \cdot 10^{-5}\,\text{m}^2\,\text{s}^{-1},$$

where we are guessing that $\kappa = \nu$ (because both κ and ν are diffusion constants). Then

$$K = \rho c_\text{p} \kappa \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1}.$$

Using this value we can estimate the heat loss on a cold day. Let's say that your skin is at $T_2 = 30\,°\text{C}$ and the air outside is $T_1 = 0\,°\text{C}$, making $\Delta T = 30\,\text{K}$. A thin T-shirt may have thickness 2 mm, so

$$F = K\frac{\Delta T}{\Delta x} \sim 0.02\,\text{W}\,\text{m}^{-1}\,\text{K}^{-1} \times \frac{30\,\text{K}}{2 \cdot 10^{-3}\,\text{m}} \sim 300\,\text{W}\,\text{m}^{-2}.$$

Damn, we want a power rather than a power per area. Ah, flux is power per area, so just multiply by a person's surface area: roughly 2 m tall and 0.5 m wide, with a front and a back. So the surface area is about $2\,\text{m}^2$. Thus, the power lost is

$$P \sim FA = 300\,\text{W}\,\text{m}^{-2} \times 2\,\text{m}^2 = 600\,\text{W}.$$

No wonder a winter day wearing only thin pants and shirt feels so cold: 600 W is large compared to human power levels. Sitting around, a person produces 100 W of heat (the basal metabolic rate). When 600 W escapes,

Which are...?

one is losing far more than the basal metabolic rate. Eventually, one's core body temperature falls. Then chemical reactions slow down. This happens for two reasons. First, almost all reactions go slower at lower temperature. Second, enzymes lose their optimized shape, so they become less efficient. Eventually you die.

One solution is jogging to generate extra heat. That solution indicates that the estimate of 600 W is plausible. Cycling hard, which generates hundreds of watts of heat, is vigorous enough exercise to keep one warm, even on a winter day in thin clothing.

Another simple solution, as parents repeat to their children: Dress warmly by putting on thick layers. Let's recalculate the power loss if you put on a fleece that is 2 cm thick. You could redo the whole calculation from scratch, but it is simpler is to notice that the thickness has gone up by a factor of 10 but nothing else changed. Because $F \propto 1/\Delta x$, the flux and the power drop by a factor of 10. So, wearing the fleece makes

$$P \sim 60\,\text{W}.$$

That heat loss is smaller than the basal metabolic rate, which indicates that one would not feel too cold. Indeed, when wearing a thick fleece, only the exposed areas (hands and face) feel cold. Those regions are exposed to the air, and are protected by only a thin layer of still air (the boundary layer). Because a large Δx means a small heat flux, the moral is (speaking as a parent): Listen to your parents and bundle up!

Only 7 Cal/minute. Not a very efficient or fun calorie burning technique.

Wait. I forgot the difference between Cal and cal.

1 Cal = 1 kcal = 1000 cal.

I love seeing these explanations (biology, and all the other things I'd never expect to see in a course 2 or 6 class) that really give us the whole story of what we're looking at.

I agree. In general, it would be nice to hear even more of this stuff although you do show us a lot. These pieces of information are what I enjoy most.

it would also be cool to comment on how long that would be

Finally an explanation in my major!! I'm no longer left out! :D

Tell it like it is, Sanjoy.

I love it!

one is losing far more than the basal metabolic rate. Eventually, one's core body temperature falls. Then chemical reactions slow down. This happens for two reasons. First, almost all reactions go slower at lower temperature. Second, enzymes lose their optimized shape, so they become less efficient. Eventually you die.

One solution is jogging to generate extra heat. That solution indicates that the estimate of 600 W is plausible. Cycling hard, which generates hundreds of watts of heat, is vigorous enough exercise to keep one warm, even on a winter day in thin clothing.

Another simple solution, as parents repeat to their children: Dress warmly by putting on thick layers. Let's recalculate the power loss if you put on a fleece that is 2 cm thick. You could redo the whole calculation from scratch, but it is simpler is to notice that the thickness has gone up by a factor of 10 but nothing else changed. Because $F \propto 1/\Delta x$, the flux and the power drop by a factor of 10. So, wearing the fleece makes

$$P \sim 60\,\mathrm{W}.$$

That heat loss is smaller than the basal metabolic rate, which indicates that one would not feel too cold. Indeed, when wearing a thick fleece, only the exposed areas (hands and face) feel cold. Those regions are exposed to the air, and are protected by only a thin layer of still air (the boundary layer). Because a large Δx means a small heat flux, the moral is (speaking as a parent): Listen to your parents and bundle up!

I've lost how we're using random walks to solve this last problem

Was just going to bring that up myself - where is the random walk here?

there wasn't one here. i mentioned this in an earlier note. we're only using the result of one of the random walks–the diffusion constant.

Yup, diffusion constant and an interesting note on why you should bundle up in the winter!

I'm confused about the connection to probability. It was much more explicit in the last sections, but how is it being used here? In a section on how to use probability, we seemed to use it briefly in the beginning with 1/2 in either direction/expected positions.

I feel like this last part of the section got slightly off track– maybe introduce that you will be doing this at the beginning of the section?

We used probabilistic analysis of a random walk to determine the Diffusion Equation and then applied this equation in cases where random walks occur (i.e. losing heat on a cold day).

yeah, the connection from probability to random walks to diffusion constant to heat loss was amazingly informative. i am impressed by how much we can figure out by just combining some premises from differing fields of study.

I actually agree with this comment. Although I like the example, I'd rather see an example that uses random walks than something it came up with (if that is even possible).

typo, delete

i really like this ... good example & well written

proportional reasoning!!!

proportional reasoning!!!

one is losing far more than the basal metabolic rate. Eventually, one's core body temperature falls. Then chemical reactions slow down. This happens for two reasons. First, almost all reactions go slower at lower temperature. Second, enzymes lose their optimized shape, so they become less efficient. Eventually you die.

One solution is jogging to generate extra heat. That solution indicates that the estimate of 600 W is plausible. Cycling hard, which generates hundreds of watts of heat, is vigorous enough exercise to keep one warm, even on a winter day in thin clothing.

Another simple solution, as parents repeat to their children: Dress warmly by putting on thick layers. Let's recalculate the power loss if you put on a fleece that is 2 cm thick. You could redo the whole calculation from scratch, but it is simpler is to notice that the thickness has gone up by a factor of 10 but nothing else changed. Because $F \propto 1/\Delta x$, the flux and the power drop by a factor of 10. So, wearing the fleece makes

$$P \sim 60\,\mathrm{W}.$$

That heat loss is smaller than the basal metabolic rate, which indicates that one would not feel too cold. Indeed, when wearing a thick fleece, only the exposed areas (hands and face) feel cold. Those regions are exposed to the air, and are protected by only a thin layer of still air (the boundary layer). Because a large Δx means a small heat flux, the moral is (speaking as a parent): Listen to your parents and bundle up!

This is cool though - it's interesting to know that the amount of power lost is inversely proportional to the thickness of my outerwear...sometimes parents do know best!

This answer is cool, but I have trouble visualizing 60W of power.

Well, you don't have to visualize that quantity. But if you want to, just think of what the text says with respect to clothing.

does your metabolic rate increase when it is cold?

That's why you shiver!

Is it a matter of thickness or density of material? There seem to be very thin, yet warm jackets these days...

Is it a matter of thickness or density of material? There seem to be very thin, yet warm jackets these days...

if your body produces 100, and you dissipate 60, then you're netting 40. do you build up heat and get too warm?

I think this is really simplistic, but at some point yes you would get warm, but not if you were just sitting there in the cold in your jacket

Not all of your body is covered. So the uncovered part would lose faster.

I have a feeling overall, although perhaps in theory you're generating enough heat, our estimations mean that you would definitely tell you were somewhere cold, not chillin on a beach in florida.

Don't take the 60 W too literally, it could easily be off by a factor of 2. But, the general point you make remains – if you are generating more heat than you lose, you'll feel warm (and probably decide to take off a jacket or find a cooling breeze).

only a parent would end a chapter with that proclamation.

great conclusion. I love it.

one is losing far more than the basal metabolic rate. Eventually, one's core body temperature falls. Then chemical reactions slow down. This happens for two reasons. First, almost all reactions go slower at lower temperature. Second, enzymes lose their optimized shape, so they become less efficient. Eventually you die.

One solution is jogging to generate extra heat. That solution indicates that the estimate of 600 W is plausible. Cycling hard, which generates hundreds of watts of heat, is vigorous enough exercise to keep one warm, even on a winter day in thin clothing.

Another simple solution, as parents repeat to their children: Dress warmly by putting on thick layers. Let's recalculate the power loss if you put on a fleece that is 2 cm thick. You could redo the whole calculation from scratch, but it is simpler is to notice that the thickness has gone up by a factor of 10 but nothing else changed. Because $F \propto 1/\Delta x$, the flux and the power drop by a factor of 10. So, wearing the fleece makes

$$P \sim 60\,\mathrm{W}.$$

That heat loss is smaller than the basal metabolic rate, which indicates that one would not feel too cold. Indeed, when wearing a thick fleece, only the exposed areas (hands and face) feel cold. Those regions are exposed to the air, and are protected by only a thin layer of still air (the boundary layer). Because a large Δx means a small heat flux, the moral is (speaking as a parent): Listen to your parents and bundle up!

great conclusion. I love it.

I feel like this conclusion might have exterior motivations, given that you are a parent and all.

Definitely ulterior motives for it as I try to convince my daughter that she needs to put on her coat when it is cold (or her sun hat when it is very sunny). The next baby is arriving in a week or so, but she'll be less able to argue, at least for a couple years.

I've also heard (and seen on television) that you can use mental "toughness" to raise your temperature enough to prevent hypothermia and death. Navy SEALS are trained to do this, I believe.

Woah, that sounds pretty awesome. I wonder what kind of physiological reactions are going on exactly?

Well you know how when you get really angry your body temperature rises..maybe they think unhappy thoughts?

Its interesting to see how 'mind over matter' actually works.

My own experience staying up all night doing physics problem sets or 24-hour take-home exams was that I had to eat a huge amount, e.g. a whole pizza, and then I was fine. I think that's because I was burning lots of calories doing so much thinking.

Consistent with that experience, I was told that Garry Kasparov, when he plays world-championship matches, eats 8000 kcal/day (instead of the usual 2500 kcal/day) and loses weight.

one is losing far more than the basal metabolic rate. Eventually, one's core body temperature falls. Then chemical reactions slow down. This happens for two reasons. First, almost all reactions go slower at lower temperature. Second, enzymes lose their optimized shape, so they become less efficient. Eventually you die.

One solution is jogging to generate extra heat. That solution indicates that the estimate of 600 W is plausible. Cycling hard, which generates hundreds of watts of heat, is vigorous enough exercise to keep one warm, even on a winter day in thin clothing.

Another simple solution, as parents repeat to their children: Dress warmly by putting on thick layers. Let's recalculate the power loss if you put on a fleece that is 2 cm thick. You could redo the whole calculation from scratch, but it is simpler is to notice that the thickness has gone up by a factor of 10 but nothing else changed. Because $F \propto 1/\Delta x$, the flux and the power drop by a factor of 10. So, wearing the fleece makes

$$P \sim 60\,\text{W}.$$

That heat loss is smaller than the basal metabolic rate, which indicates that one would not feel too cold. Indeed, when wearing a thick fleece, only the exposed areas (hands and face) feel cold. Those regions are exposed to the air, and are protected by only a thin layer of still air (the boundary layer). Because a large Δx means a small heat flux, the moral is (speaking as a parent): Listen to your parents and bundle up!

I don't know if this is coming, but it'd be great to relate random walks to errors in estimation. If we are using divide and conquer, our log(error factor) is a variation of the random walk discussed here, where each step itself has some (unknown but approximable, e.g. by the log normal curve) distribution.

I think this would tie together several different topics, and would reinforce the point that our error's variance grows (linearly?) with the number of divisions we make, and therefore the standard deviation grows at a lesser rate.

Yes, that would work well. For error estimations, it's still a one-dimensional walk. But the step sizes are now not fixed; instead, each factor contributes one step, and its size is normally distributed with mean 0 and standard deviation equal to the number of log units of plus/minus (e.g. if it's 10^(a +/- b) then the standard deviation is b).

How about using the random walk analysis to use Brownian motion to figure out the size of molecules?

That would be an interesting section and would also allow you to firm up your 3 cubic angstrom assertion.

That would be interesting to see. I'd like to see that worked out.

9
Springs

9.1 Why planets orbit in ellipses 175
9.2 Musical tones 181
9.3 Waves 186
9.4 Precession of planetary orbits 227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

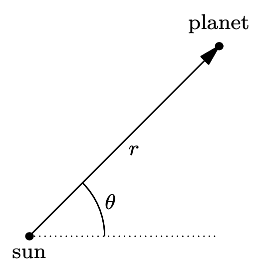

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

Global comments

There is way too much mathematical explanations in this section for me to fully grasp the subject. It's very non-intuitive for me and I am having trouble reading through the section.

9
Springs

9.1 Why planets orbit in ellipses 175
9.2 Musical tones 181
9.3 Waves 186
9.4 Precession of planetary orbits 227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

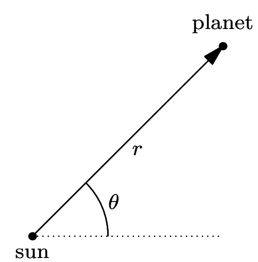

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

Do you really mean contains, or is it more, "can be approximated as"

When I read the title, I was thinking, "there's no way this is going to only be about springs..."

really? " every "physical process?

yes. almost every physical process. it's not really that it literally contains a spring, just that they all act similarly ... the spring math can be used just about everywhere.

really, they "can be approximated as" a spring...

Contains or can be modeled as a spring?

I think he's treating those expressions as equivalent.

Ooh this promises to be interesting...

yeah, great opening sentnence

This is surprising. I look forward to reading more!

Yeah...im very curious to see how this is going to make sense

This sounds intriguing!

Sounds very 8.01!

Read this section for Monday (memo due Monday at 9am). It starts our final tool: the spring approximation.

Perhaps you could put a couple lines in the beginning to outline how we will go about addressing the topic (springs) with the example as you've done in previous readings. For me it makes the analysis easier to follow.

I really like this idea...i think it would have helped me get my head around the concept _a lot_ faster.

is there a reason that this isn't m[sun] >> m[planet]?

9
Springs

9.1 Why planets orbit in ellipses 175
9.2 Musical tones 181
9.3 Waves 186
9.4 Precession of planetary orbits 227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \quad (9.1)$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \quad (9.2)$$

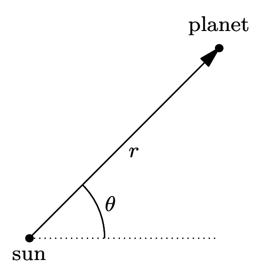

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

is there a reason that this isn't m[sun] >> m[planet]

is there a reason that this isn't m[sun] >> m[planet]?

is there a reason that this isn't m[sun] >> m[planet]?

why? do we just take the gravitational constant for granted for that distance?

Would also like an answer to this. Also, it's a little unclear if it's the sun that's infinitely massive or the planet. Obviously it has to be the sun, but my first reaction was to think it's the planet and I was a little confused.

wouldn't the sun be moving around it?

He is referring to the sun here, not the planet

Is this a fair assumption? I mean, sure, the sun is VASTLY greater than a planet but if you compare say our sun and Jupiter, the difference is completely within several orders of magnitude and not near infinite.

the reason for the assumption is that we're assuming the sun stays in one place. if it weren't assumed infinitely massive, then we'd have to consider the (much smaller) orbit of the sun as well.

9
Springs

9.1 Why planets orbit in ellipses 175
9.2 Musical tones 181
9.3 Waves 186
9.4 Precession of planetary orbits 227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

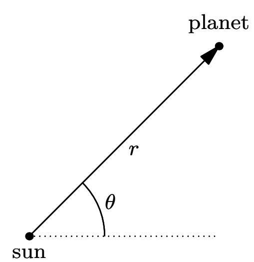

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

Why do we have this condition?

I think because otherwise the planet would exert a gravitational force on the sun and it would mess up our calculations.

either the gravational force idea, or that we just want to assume the center of mass lies roughly inside the sun

Without assuming this, they would both essentially orbit around some point in space (like a binary star system). Making this assuming basically assumes that the sun doesn't move, and the center of the orbiting system lies at the center of the sun.

Additionally, the mass of the sun (or any star) is so many more times greater than the mass of its orbiting planet that this approximation is very accurate and turns out to basically make the math easier.

I agree. The sun's mass is so much larger that it makes the effects of the earth on the sun negligible and therefore you shouldn't have to calculate them

After reading this section it looks like we need to explain why we assume that the sun is infinitely massive as well as under what circumstances that assumption should be made and that it won't affect the equations you give later. There seems to be a lot of confusion in this section about this term.

As mentioned above by 08:18, the correct answer should be that an infinite mass will not move when orbited by another finite mass.

When people say it makes the "effect on the Earth" neglible, I hope you're NOT referring to force. Forces are equal and opposite in nature. You might be tempted to think that the sun exerts a larger force on the Earth, but there is no such thing as an object exerting its own force on another object. Forces are not possessive, but are the *NET* effect of pairs of objects interacting with each other. Therefore, each object in that pair will feel the same force, in opposite directions.

I understand why an assumption like this must be made to simplify the equations. however, is there a reason that the sun is "infinitely" massive as opposed to it's mass being 'much greater than' that of the planet, ie m[planet] << m[sun]

9
Springs

9.1 Why planets orbit in ellipses 175
9.2 Musical tones 181
9.3 Waves 186
9.4 Precession of planetary orbits 227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

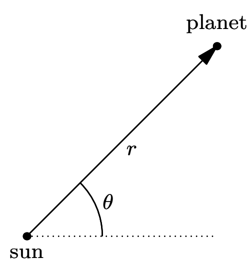

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

After reading all the comments, I see that I stated it badly and partly incorrectly (if the sun were actually infinitely massive, then the planet would orbit at infinite speed!). A better statement is that the mass of the planet is very small compared to the mass of the sun (in E&M, that's the "test-charge condition"). Then the sun hardly moves, and the kinetic energy of the sun is small compared to the kinetic energy of the planet.

Here's a scaling argument to show that. The KE (kinetic energy) of the sun is roughly M*V^2 where M = mas sof sun and V =speed of sun. Conservation of momentum says M*V=m*v (where m=mass of planet, v = speed of planet). So the ratio of kinetic energies, M*V^2/(m*v^2), is m/M (after a few steps of algebra). So what I really mean is : m/M << 1.

How do you assume something as infinitely massive and then use that assumption to get energy per mass?

I think the suns mass is assume to be infinite, not the planets.

another assumption that is made but i havent been able to find so far is that the sun and planet are being treated as point masses.

is there a reason that this isn't m[sun] >> m[planet]?

why is this negative?

9 Springs

9.1 Why planets orbit in ellipses 175
9.2 Musical tones 181
9.3 Waves 186
9.4 Precession of planetary orbits 227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

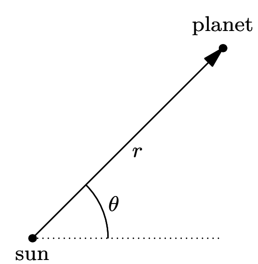

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

I'm not familiar with planetary physics and the associated formulas. Where exactly does this come from? Perhaps defining E before you divide by m would make it more clear.

This is just the gravitational potential energy (-GMm/r) plus the kinetic energy of the plane moving about the sun (1/2 *m*v^2), then divided by m (the mass of the planet). M is the mass of the sun, and G is newton's Gravitation constant, and r is the distance from the sun.

This is a pretty standard representation of Energy (Just PE + KE), just in the planetary sense. This is fairly common for my major, but it may help if there was a basic explanation of how this is the potential + kinetic energies

I agree that its pretty recognizable but a quick debrief about energy at the intro of the section wouldn't be so bad.

Or maybe just starting a little less derived would help. As in showing where this equation came from would help clarify matters.

Even just below each quantity, just labeling it. Like "energy/mass = gravitational energy + kinetic energy"

So how does the sun's infinite mass affect this equation? since there are two masses, I'm assuming that one is the planet and one is the sun

I'm also not sure about this, maybe that assumption is what keeps the sun from moving as well as the earth, which would create a more dynamic system.

Yeah, the infinity makes things a bit strange here. The first term in E/m contains M, the sun's mass, so it should go to infinity.

oh so were going to say that E/m goes to zero so that we can solve the problem easier... it would be nice to see how you arrived at this

I'm not sure that we're going to have the energy go to zero. Also, this equation comes from the fact that the energy of the earth is the sum of the gravitational potential energy and the kinetic energy.

^ Ditto.

using the real mass and not the infinite assumption of course. correct?

9
Springs

9.1	Why planets orbit in ellipses	175
9.2	Musical tones	181
9.3	Waves	186
9.4	Precession of planetary orbits	227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

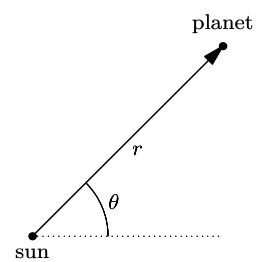

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

If the sun's mass is assumed to be massive, why is the sun's mass in the equation. If we truly are making the assumption, wouldn't the equation go to -infinity?

Because the potential energy (essentially how much its gravity affects what is orbiting around it) is dependent upon the mass of that body.(Larger mass, more gravity) The sun is very massive compared to that of Earth, but it definitely isn't infinity.

But the line directly before this says "the sun (assumed to be infinitely massive)." That seems pretty straightforward to me. But how does that apply to this? A previous comment said something about center of mass, could it be just to make that assumption and nothing else?

I'm not sure why we made the original assumption, but it would be nice to have it spelled out. I had the same question here.

In order for this theory to be valid, it would require an infinite mass. However the equation used requires a real mass. I understand, but it would be more clear with further explanation.

would the theory still be valid if "m[sun] >> m[planet]"? I feel like this statement still allows us to make the assumption that the center of mass of the system is the center of mass of the sun and still accept that GM/r is not infinity.

Tangential velocity, I assume, not angular. Not that there really would be confusion at this point, but it might be worth pointing out.

I love polar coordinates for spherical motion (go figure...). In high school, our teacher wanted us to do circular motion in Cartesian coordinates "for instructive purposes and to make you appreciate polar coordinates."

I agree, it makes it much easier to understand a problem like this when it's in polar coordinates.

Wish my teacher did that, maybe I would like polar coordinate more.

my only problem with polar coordinates was that we were always forced to switch between them and cartesian

9
Springs

9.1 Why planets orbit in ellipses 175
9.2 Musical tones 181
9.3 Waves 186
9.4 Precession of planetary orbits 227

Almost every physical process contains a spring! The first example of that principle shows a surprising place for a spring: planetary orbits.

9.1 Why planets orbit in ellipses

For a planet moving around the sun (assumed to be infinitely massive), the planet's energy per mass is

$$\frac{E}{m} = -\frac{GM}{r} + \frac{1}{2}v^2, \tag{9.1}$$

where m is the mass of the planet, G is Newton's constant, M is the mass of the sun, r is the planet's distance from the sun, and v is the planet's speed.

In polar coordinates, the kinetic energy per mass is

$$\frac{1}{2}v^2 = \frac{1}{2}\left[r^2\left(\frac{d\theta}{dt}\right)^2 + \left(\frac{dr}{dt}\right)^2\right]. \tag{9.2}$$

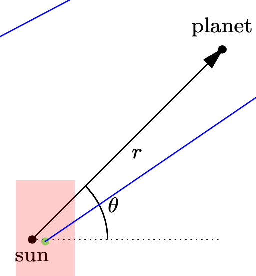

This energy contains two coordinates, which is one too many for an easy solution. We can get rid of a coordinate by incorporating a conservation law,

i'm glad you distinguish between the potential and kinetic energy terms in the above equation.

agreed, but I feel like it should happen with the equation ... say in words before in the equation ... it would have helped me remember why this eq sooner.

how do you convert to polar coordinates? I forgot

Should this be an equation most of us should know? becasue i dont't think I've ever seen it before

It's a transformation of the kinetic energy term in equation 9.2 to polar coordinates in 2 dimensions. dtheta/dt is the change in theta over time, which is a velocity. dr/dt is the change in radius over time, which is another velocity.

Interesting. Thank you for the explanation, although the next page also helps with the understanding.

Thanks, I was staring at this wondering what it was. I wonder if the general audience needs to be reminded of this?

even though I am aware of this equation...I think you should break it down and explain in greater detail

So, mass is missing in this equation because it's infinitely massive?

Because this equation is for the 'kinetic energy per mass'. The quantity E/m.

Wouldn't we be losing information by doing this? Why would we want to get rid of a coordinate?

technically, we're not just throwing it out. we're "getting rid of it" by canceling it out with another known equation. there is no loss.

Does the fact that the sun moves in an orbit itself complicate this significantly?

angular momentum. The angular momentum per mass is

$$\ell = r^2 \frac{d\theta}{dt}, \tag{9.3}$$

The angular momentum per mass ℓ allow us to eliminate the θ coordinate by rewriting the $\frac{d\theta}{dt}$ term:

$$r^2 \left(\frac{d\theta}{dt}\right)^2 = \frac{\ell^2}{r^2}. \tag{9.4}$$

Therefore,

$$\frac{1}{2}v^2 = \frac{1}{2}\left[\frac{\ell^2}{r^2} + \left(\frac{dr}{dt}\right)^2\right]; \tag{9.5}$$

and

$$\frac{E}{m} = \underbrace{-\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}}_{V_{\text{eff}}} + \underbrace{\frac{1}{2}\left(\frac{dr}{dt}\right)^2}_{\text{KE per mass}}. \tag{9.6}$$

Because the gravitational force is central (toward the sun), the planet's angular momentum, when computed about the sun, is constant. Therefore, the only variable in E/m is r. This energy per mass describes the motion of a particle in one dimension (r). The first two terms are the potential (the potential energy per mass); they are called the effective potential V_{eff}. The final term is the particle's kinetic energy per mass.

Now let's study the just the effective potential:

$$V_{\text{eff}} = -\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}. \tag{9.7}$$

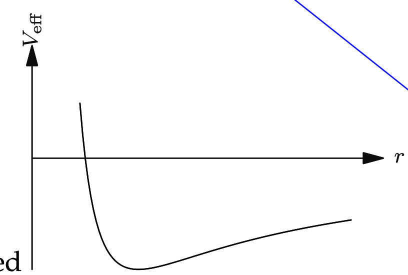

The first term is the actual gravitational potential; the second term, which originated from the tangential motion, is called the centrifugal potential. To understand how they work together, let's make a sketch. For almost any sketch, the first tool to pull out is easy cases. Here, the easy cases are small and large r. At small r, the centrifugal potential is the important term because its

angular momentum isn't really a conservation law. There is a conservation law for it, but this makes it sound like you're saying the term itself is the law. It'd be like saying "energy is a conservation law".

Why r squared?

I would also like an answer to this question.

I think it's because of the definition of angular momentum L=r x p, where "r" is the moment arm and "p" is the linear momentum (which is mass x velocity). Since we only need tangential velocity for linear momentum, we use r*d_theta.

Of course we divide all this by mass since we want angular momentum per mass.

This makes a lot of sense. I was having trouble with the units for a second.

is this the symbol for angular momentum?it looks like an l

How does this help us if we don't know what the angular momentum is? We are basically substituting one variable for another

It helps because the angular momentum is constant, so you replace a changing quantity (d(theta)/dt) with a constant. it's the maxim from the invariance chapter: "When there is change, look for what does not change."

A constant, even an unknown one, is much easier to handle than a changing quantity. (For example, the constant has zero derivative.)

should be: allows

Hmm, quite clever. It took me a second to realize how we went from 9.3 to 9.4, but I see we just squared both sides and divided by r^2.

I agree. These simple substitutions are what always trips me up, because I am on the look out for something more complicated.

Ha, neat!

angular momentum. The angular momentum per mass is

$$\ell = r^2 \frac{d\theta}{dt}, \tag{9.3}$$

The angular momentum per mass ℓ allow us to eliminate the θ coordinate by rewriting the $\frac{d\theta}{dt}$ term:

$$r^2\left(\frac{d\theta}{dt}\right)^2 = \frac{\ell^2}{r^2}. \tag{9.4}$$

Therefore,

$$\frac{1}{2}v^2 = \frac{1}{2}\left[\frac{\ell^2}{r^2} + \left(\frac{dr}{dt}\right)^2\right]; \tag{9.5}$$

and

$$\frac{E}{m} = \underbrace{-\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}}_{V_{\text{eff}}} + \underbrace{\frac{1}{2}\left(\frac{dr}{dt}\right)^2}_{\text{KE per mass}}. \tag{9.6}$$

Because the gravitational force is central (toward the sun), the planet's angular momentum, when computed about the sun, is constant. Therefore, the only variable in E/m is r. This energy per mass describes the motion of a particle in one dimension (r). The first two terms are the potential (the potential energy per mass); they are called the effective potential V_{eff}. The final term is the particle's kinetic energy per mass.

Now let's study the just the effective potential:

$$V_{\text{eff}} = -\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}. \tag{9.7}$$

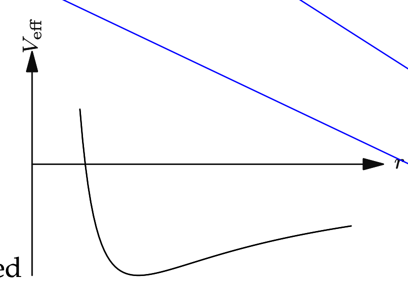

The first term is the actual gravitational potential; the second term, which originated from the tangential motion, is called the centrifugal potential. To understand how they work together, let's make a sketch. For almost any sketch, the first tool to pull out is easy cases. Here, the easy cases are small and large r. At small r, the centrifugal potential is the important term because its

does the radius change? it shouldn't so shouldn't dr/dt be zero?

I think that the point of this section is to explain why r does change. An ellipse is like a circle with varying radius.

I feel like it's a vote of confidence for my gut that I forgot to do this reading before lecture (oops! sorry!) and I still got the lecture demo right by intuition.

I guess I'm getting confused as to what velocity we're referring to. Above dTheta/dt was velocity in the momentum equation. But here dr/dt is velocity in the kinetic energy term.

If you look at the 9.2 equation, you will see that both dTheta/dt and dr/dt are components of velocity in the kinetic energy equation.

Took me a while to draw that conclusion. Lots to remember in a short period of time!

Be careful! dTheta/dt is an angular velocity, the actual component of velocity is r*dTheta/dt/

what is Veff?

I like it when you break down what the terms in the equation translate too

I agree, I think it could definitely help with the initial energy equation introduced on the first page as well.

I'm a big fan of how it has the little subtitles and the brackets throughout the chapters.

Oh ok,s o that's why we substituted earlier for a seemingly unknown L. However, I'm not sure if I would have fully understood this without the comment right here.

Eliminating the dependency on its angle by using angular momentum, and then realizing angular momentum is constant is a clever trick. I don't think I would think about making this move if I wasn't shown how to do it.

because they're perpendicular, right?

wait what? i don't follow....

Is there a particular reason they have that name?

angular momentum. The angular momentum per mass is

$$\ell = r^2 \frac{d\theta}{dt}, \tag{9.3}$$

The angular momentum per mass ℓ allow us to eliminate the θ coordinate by rewriting the $\frac{d\theta}{dt}$ term:

$$r^2 \left(\frac{d\theta}{dt}\right)^2 = \frac{\ell^2}{r^2}. \tag{9.4}$$

Therefore,

$$\frac{1}{2}v^2 = \frac{1}{2}\left[\frac{\ell^2}{r^2} + \left(\frac{dr}{dt}\right)^2\right]; \tag{9.5}$$

and

$$\frac{E}{m} = \underbrace{-\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}}_{V_{\text{eff}}} + \underbrace{\frac{1}{2}\left(\frac{dr}{dt}\right)^2}_{\text{KE per mass}}. \tag{9.6}$$

Because the gravitational force is central (toward the sun), the planet's angular momentum, when computed about the sun, is constant. Therefore, the only variable in E/m is r. This energy per mass describes the motion of a particle in one dimension (r). The first two terms are the potential (the potential energy per mass); they are called the effective potential V_{eff}. The final term is the particle's kinetic energy per mass.

Now let's study the just the effective potential:

$$V_{\text{eff}} = -\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}. \tag{9.7}$$

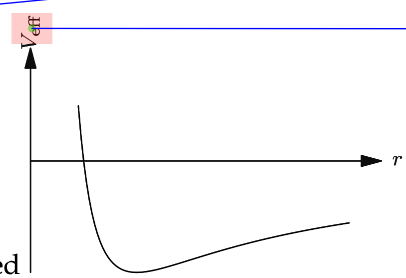

The first term is the actual gravitational potential; the second term, which originated from the tangential motion, is called the centrifugal potential. To understand how they work together, let's make a sketch. For almost any sketch, the first tool to pull out is easy cases. Here, the easy cases are small and large r. At small r, the centrifugal potential is the important term because its

We need an alphabet with more letters... When I first saw Veff in the equation I thought it had something to do with velocity since the other term was kinetic energy.

Yeah I was confused by that notion initially as well. It would be nice to see either a different notion or to see this current one (V_eff) defined earlier.

Right, I always thought potential was U.

I think you're think of potential energy – which this is not, it's potential energy per mass. It's similar but different enough that the V reminds us that we're not dealing with units of J but J/kg.

at least it's capitalized.

What is the difference between regular potential and effective potential. Does one take into account the mass of both of the planet and the sun?

Right here I started to wonder where springs or estimation were going to come in at all. Of course, I'll keep reading, it but it would be kind of nice to have had a preview before here. So far it's read like a physics textbook.

define effective?? I know you described this in the previous paragraph, but what deliniates it from other potential measurements?

It's different from other potential measurements because it's not simply gravitational/electrical/chem potential energy. What it is, however, is a measure of energy which depends only on distance (not, for example, on speed), so it *behaves* like a potential energy.

I think this is a typo?

I suspect that it may have been difficult to label the graph this way, but I think it's actual perfectly clear the if you rotate it 90 degrees to the right.

Why would we cant to rotate the axis? We are looking at the potential as a function of radius not the other way around. I think it would be much more confusing if we flipped it

not rotate the axis, rotate the axis label. and i agree that it would have been a lit easier to read without turning my head or the page

I agree, and I tried to rotate them. But I couldn't convince the figure-drawing program ("asymptote") to do it, but I'll try again for the next version.

angular momentum. The angular momentum per mass is

$$\ell = r^2 \frac{d\theta}{dt}, \tag{9.3}$$

The angular momentum per mass ℓ allow us to eliminate the θ coordinate by rewriting the $\frac{d\theta}{dt}$ term:

$$r^2 \left(\frac{d\theta}{dt}\right)^2 = \frac{\ell^2}{r^2}. \tag{9.4}$$

Therefore,

$$\frac{1}{2}v^2 = \frac{1}{2}\left[\frac{\ell^2}{r^2} + \left(\frac{dr}{dt}\right)^2\right]; \tag{9.5}$$

and

$$\frac{E}{m} = \underbrace{-\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}}_{V_{\text{eff}}} + \underbrace{\frac{1}{2}\left(\frac{dr}{dt}\right)^2}_{\text{KE per mass}}. \tag{9.6}$$

Because the gravitational force is central (toward the sun), the planet's angular momentum, when computed about the sun, is constant. Therefore, the only variable in E/m is r. This energy per mass describes the motion of a particle in one dimension (r). The first two terms are the potential (the potential energy per mass); they are called the effective potential V_{eff}. The final term is the particle's kinetic energy per mass.

Now let's study the just the effective potential:

$$V_{\text{eff}} = -\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}. \tag{9.7}$$

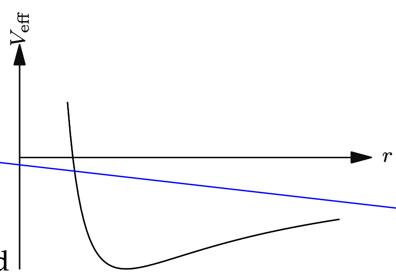

The first term is the actual gravitational potential; the second term, which originated from the tangential motion, is called the centrifugal potential. To understand how they work together, let's make a sketch. For almost any sketch, the first tool to pull out is easy cases. Here, the easy cases are small and large r. At small r, the centrifugal potential is the important term because its

I think graphs are always helpful to help with our understanding of the material.

You might want to explain intuitively what "effective potential" means; otherwise, it's just another term being thrown out.

Since I have no knowledge as to what "effective potential" is referring to and how it differs from a normal "potential", I agree.

"effective" just means we take the normal potential and adjust it somehow, so it seems like its potential is something else.

Agreed. For all practical purposes you can just consider the effective potential to be like any other potential. The only difference is that the effective potential includes some other quantity (angular momentum here) that isn't necessarily a "potential" as we might usually think of one. It's really just a way of lumping the terms in our equation to make it easier to handle.

I agree with the fourth person, there's really not much to the term effective and I think for this case it's explained in the line above so it isn't too confusing.

I think it would be nice just to have a short sentence that explains what effective potential is. It would be a nice refresher for people who already know and a helpful piece of information for people who are unfamiliar.. plus it wouldn't take up too much space.

I don't see why we need the term effective potential at all- right after we lumped everything together and called it the "effective potential", we take it apart again and analyze its parts separately anyway

alternatively, explain what the potential is and how the effective potential differs. It'd be nice to see the base case here.

I agree. I dont know why we really need to rename it. I understood the first time that you wrote this equation above, that it consisted of the potential and kinetic energy components. If you are going to name this term or clarify its components for readers, I would recommend doing it earlier in the reading.

If we're assuming the sun is infinitely massive, why doesn't this term drop out or just make the expression infinite?

angular momentum. The angular momentum per mass is

$$\ell = r^2 \frac{d\theta}{dt}, \tag{9.3}$$

The angular momentum per mass ℓ allow us to eliminate the θ coordinate by rewriting the $\frac{d\theta}{dt}$ term:

$$r^2 \left(\frac{d\theta}{dt}\right)^2 = \frac{\ell^2}{r^2}. \tag{9.4}$$

Therefore,

$$\frac{1}{2}v^2 = \frac{1}{2}\left[\frac{\ell^2}{r^2} + \left(\frac{dr}{dt}\right)^2\right]; \tag{9.5}$$

and

$$\frac{E}{m} = \underbrace{-\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}}_{V_{\text{eff}}} + \underbrace{\frac{1}{2}\left(\frac{dr}{dt}\right)^2}_{\text{KE per mass}}. \tag{9.6}$$

Because the gravitational force is central (toward the sun), the planet's angular momentum, when computed about the sun, is constant. Therefore, the only variable in E/m is r. This energy per mass describes the motion of a particle in one dimension (r). The first two terms are the potential (the potential energy per mass); they are called the effective potential V_{eff}. The final term is the particle's kinetic energy per mass.

Now let's study the just the effective potential:

$$V_{\text{eff}} = -\frac{GM}{r} + \frac{1}{2}\frac{\ell^2}{r^2}. \tag{9.7}$$

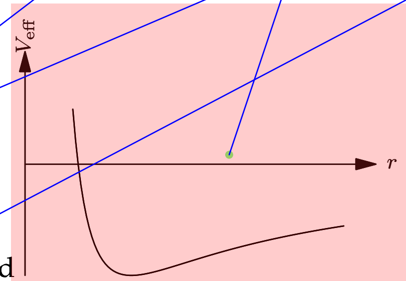

The first term is the actual gravitational potential; the second term, which originated from the tangential motion, is called the centrifugal potential. To understand how they work together, let's make a sketch. For almost any sketch, the first tool to pull out is easy cases. Here, the easy cases are small and large r. At small r, the centrifugal potential is the important term because its

I like this sketch.

This looks like a lot of other things we've seen in class. It might be interesting at the end to figure otu which graphs look like they relate to each other and find out why....nature seems to make things work in certain ways, and just keeps repeating them!

I definitely agree with this I think it would be really cool to look at relations that have come out of the work we've done in this class that we might never see otherwise.

How was this graph compiled? Easy cases?

I'm curious about this too - was it compiled with easy cases or was it simply plotted with software?

Actually, this is exactly the same plot we saw when estimating the size of an atom using electrostatic potential and Heisenberg uncertainty.

as opposed to the fake potential?

I think it's the "actual" because we're used to thinking about potential as gravitational potential. The other term isn't part of the "actual" gravitational potential that we usually see.

I think I see where this is going already, and it's very cool!

these potentials are kinda confusing

I always like to see a familiar technique come into play for a new problem!!

I like to see when easy cases are used. I think easy cases are the basis for solving a lot of these problems and it almost could have been taught earlier.

I agree–easy cases are easily applied since everything can have easy cases, and the nature of easy cases is quite intuitive as well, so it's always a good starting point.

yeah I agree.. I feel like easy cases should have been the first or second topic, since it's a simply technique that most people use without realizing it.

$1/r^2$ overwhelms the $1/r$ in the gravitational potential. At large r, the gravitational potential is the important term, because its $1/r$ approaches zero more slowly than does the $1/r^2$ in the centrifugal potential. Therefore,

$$V_{\text{eff}} \propto \begin{cases} 1/r^2 & \text{(small } r) \\ -1/r & \text{(large } r) \end{cases} \tag{9.8}$$

If this analysis and sketch look familiar, that's because they are. The effective potential has the same form as the energy in hydrogen (Section 7.3 or `r26-lumping-hydrogen.pdf`); therefore, our conclusions about planetary motion will generalize to hydrogen.

Imagine the planet orbiting in a circular orbit. In that orbit, r remains constant so dr/dt and d^2r/dt^2 are both zero. For that to be true, the particle must live where the effective potential V_{eff} is flat – in other words, at its minimum at $r = r_0$.

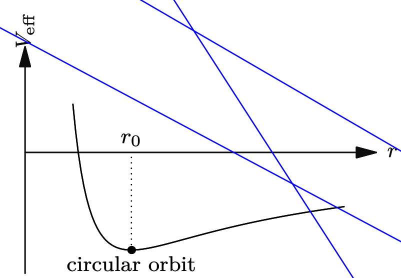

Now perturb the orbit by kicking the planet slightly outwards. That kick does not change the angular momentum ℓ, because angular momentum depends on the tangential velocity. But it gives the planet a nonzero radial velocity ($dr/dt \neq 0$). Thus, it now has r-coordinate kinetic energy (the θ-coordinate kinetic energy is taken care of by the centrifugal-potential piece of the effective potential). The orbital radius r then varies between the extremes where the r kinetic energy turns completely into effective-potential energy. Those are the two points where the horizontal line intersects the effective potential, and the corresponding radii are the minimum and maximum orbital distances.

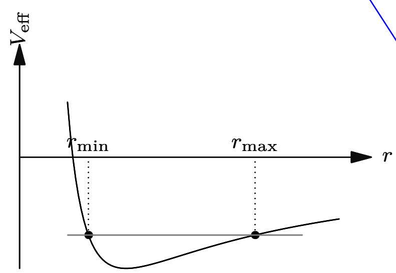

Comments on page 3

This paragraph is a great and concise explanation.

It also makes a lot of intuitive sense in terms of how planetary orbits work. I like this paragraph and the diagram a lot.

This is one of my favorite quick sketch techniques that we use in this class that has been incredibly useful in my other classes

I agree, I love this strategy.

I agree. I've found this very useful as well.

same...this is really helpful!

Yeah I've also been using this for my other classes.

I agree as well, I had a feeling that this model would emulate the other hydrogen curve from when we first saw the equation.

This is asymptotic approximation. right?

it's like this kind of graph is everywhere, especially when we calculate energy!

I like that you point this out because I'm not sure I would have noticed this myself and it's a really interesting piece of information!

Me too, I especially like the reference in the next line to the exact pages we can find it.

This sort of thing has popped up so many times in my physics courses.

agreed. i like the very specific reference instead of something general

agreed. i like the very specific reference instead of something general

quantum physics is so interesting...very small and very large objects follow similar rules

$1/r^2$ overwhelms the $1/r$ in the gravitational potential. At large r, the gravitational potential is the important term, because its $1/r$ approaches zero more slowly than does the $1/r^2$ in the centrifugal potential. Therefore,

$$V_{\text{eff}} \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{9.8}$$

If this analysis and sketch look familiar, that's because they are. The effective potential has the same form as the energy in hydrogen (Section 7.3 or `r26-lumping-hydrogen.pdf`); therefore, our conclusions about planetary motion will generalize to hydrogen.

Imagine the planet orbiting in a circular orbit. In that orbit, r remains constant so dr/dt and d^2r/dt^2 are both zero. For that to be true, the particle must live where the effective potential V_{eff} is flat – in other words, at its minimum at $r = r_0$.

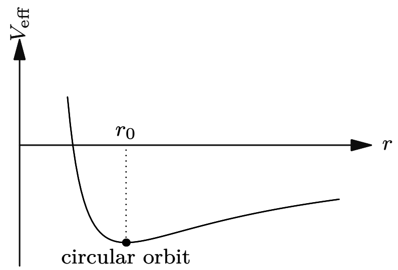

Now perturb the orbit by kicking the planet slightly outwards. That kick does not change the angular momentum ℓ, because angular momentum depends on the tangential velocity. But it gives the planet a nonzero radial velocity ($dr/dt \neq 0$). Thus, it now has r-coordinate kinetic energy (the θ-coordinate kinetic energy is taken care of by the centrifugal-potential piece of the effective potential). The orbital radius r then varies between the extremes where the r kinetic energy turns completely into effective-potential energy. Those are the two points where the horizontal line intersects the effective potential, and the corresponding radii are the minimum and maximum orbital distances.

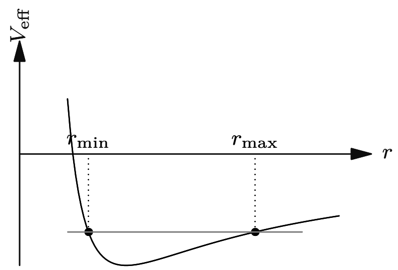

I like how this connected the readings, and how there's a direct reference to it, making it easier to find.

yeah I agree...I feel like in earlier sections when you wanted to reference a previous example you just mentioned what part of the book it was from, which was not as helpful. Instead here you give the exact section number AND the pdf which is really useful in the context of our class!

thinking along the lines of 1:22 ... when you publish the book, i think that including page numbers with your section references would be very helpful

This is pretty cool - but it makes sense when you just think about how similar the atom and planetary models are.

That's true, but it's still surprising given that at the atomic scale we have to worry about quantum effects and at the planetary scale quantum definitely does not apply.

I was thinking about the reason for the similarity a bit more. In both cases it's due to angular momentum. For the planet, the 1/r^2 term arises from using angular-momentum conservation to replace the d(theta)/dt term. For hydrogen, the 1/r^2 term arises from the uncertainty principle (Delta-p)*(Delta-x) = hbar. And hbar is the quantum of angular momentum (the size of a packet of angular momentum).

It does, which is weird when you think about how big planets are in relation to atoms but the physics are the same for the two situations, but then apparently when you go from atoms to the quantum level the physics changes.

$1/r^2$ overwhelms the $1/r$ in the gravitational potential. At large r, the gravitational potential is the important term, because its $1/r$ approaches zero more slowly than does the $1/r^2$ in the centrifugal potential. Therefore,

$$V_{\text{eff}} \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{9.8}$$

If this analysis and sketch look familiar, that's because they are. The effective potential has the same form as the energy in hydrogen (Section 7.3 or `r26-lumping-hydrogen.pdf`); therefore, our conclusions about planetary motion will generalize to hydrogen.

Imagine the planet orbiting in a circular orbit. In that orbit, r remains constant so dr/dt and d^2r/dt^2 are both zero. For that to be true, the particle must live where the effective potential V_{eff} is flat – in other words, at its minimum at $r = r_0$.

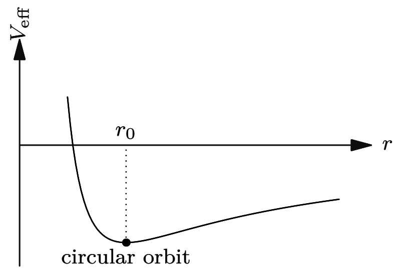

Now perturb the orbit by kicking the planet slightly outwards. That kick does not change the angular momentum ℓ, because angular momentum depends on the tangential velocity. But it gives the planet a nonzero radial velocity ($dr/dt \neq 0$). Thus, it now has r-coordinate kinetic energy (the θ-coordinate kinetic energy is taken care of by the centrifugal-potential piece of the effective potential). The orbital radius r then varies between the extremes where the r kinetic energy turns completely into effective-potential energy. Those are the two points where the horizontal line intersects the effective potential, and the corresponding radii are the minimum and maximum orbital distances.

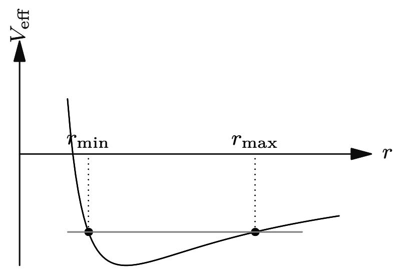

Wow, I am surprised by this conclusion.

Exactly! similar charts. It also reminds me a bit of what the graphs from the plane drag looked like when we added them to find the ideal height or velocity of the plane

Again I'd love to maybe spend the last day or two looking at the similarities we've derived in different areas of the physical world, it's crazy how much behaves in a certain way

Is this saying that planets orbit the sun in a similar way that an electron "orbits" the nucleus of a hydrogen atom? If not, in what way does this problem generalize to hydrogen?

For hydrogen, you also would make an effective potential The difference is that we can't treat hydrogen classically because it is so small.

The coloumb potential and the gravitational potential have the same form though, so it's not surprising that the graph is the same.

Yeah, this is very interesting. At first it seems quite intuitive that these two models would be similar, but then when you think about it, it's pretty crazy that things on the scale of hydrogen and things on the scale of planets, which use obey laws and forces, still can be generalized

For the last day, I've planned (following the tradition I learnt from Donald Knuth's courses) a short summary and then mostly "your turn", i.e. ask any question and I'll do my best. Explaining similarities would be a fun question.

I love how we can connect things from opposite ends of the spectrum- hydrogen atoms and planets behaving in the same way

$1/r^2$ overwhelms the $1/r$ in the gravitational potential. At large r, the gravitational potential is the important term, because its $1/r$ approaches zero more slowly than does the $1/r^2$ in the centrifugal potential. Therefore,

$$V_{\text{eff}} \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{9.8}$$

If this analysis and sketch look familiar, that's because they are. The effective potential has the same form as the energy in hydrogen (Section 7.3 or `r26-lumping-hydrogen.pdf`); therefore, our conclusions about planetary motion will generalize to hydrogen.

Imagine the planet orbiting in a circular orbit. In that orbit, r remains constant so dr/dt and d^2r/dt^2 are both zero. For that to be true, the particle must live where the effective potential V_{eff} is flat – in other words, at its minimum at $r = r_0$.

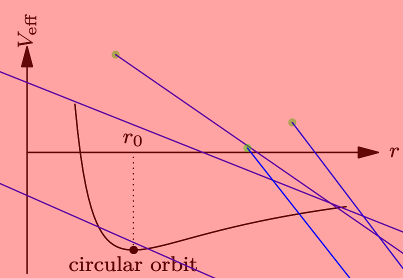

Now perturb the orbit by kicking the planet slightly outwards. That kick does not change the angular momentum ℓ, because angular momentum depends on the tangential velocity. But it gives the planet a nonzero radial velocity ($dr/dt \neq 0$). Thus, it now has r-coordinate kinetic energy (the θ-coordinate kinetic energy is taken care of by the centrifugal-potential piece of the effective potential). The orbital radius r then varies between the extremes where the r kinetic energy turns completely into effective-potential energy. Those are the two points where the horizontal line intersects the effective potential, and the corresponding radii are the minimum and maximum orbital distances.

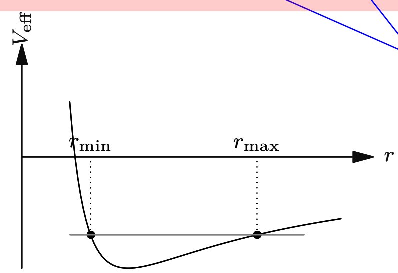

In response to a similar observation in another thread, I posted the following (which I'll just repeat here):

I was thinking about the reason for the similarity a bit more. In both cases it's due to angular momentum. For the planet, the 1/r^2 term arises from using angular-momentum conservation to replace the d(theta)/dt term. For hydrogen, the 1/r^2 term arises from the uncertainty principle (Delta-p)*(Delta-x) = hbar. And hbar is the quantum of angular momentum (the size of a packet of angular momentum).

[My double posting of the response reminds me of a NB project that could make a good UROP (or I might work on it myself at Olin next year): discussions that can be merged, modified, and otherwise improved for the benefit of all the readers.

A model for that is stackoverflow.com, which I learnt about from Micah Siegel (my roomate in graduate school, whose college roommate created stackoverflow). In the FAQ at http://stackoverflow.com/faq look at the section on "Reputation" for an interesting list of the possibilities that could happen if students could be NB moderators.]

so we are trying to minimize the energy, just like the hydrogen

ok, sweet, my previous comment has been answered

It's nice seeing the results visualized this way.

It would be interesting to see a graph of the circularity of the orbit compared to r.

Hmm, I guess that's true, but it's not really relevant here. You can look that up whenever, this is specifically about elliptical orbits.

$1/r^2$ overwhelms the $1/r$ in the gravitational potential. At large r, the gravitational potential is the important term, because its $1/r$ approaches zero more slowly than does the $1/r^2$ in the centrifugal potential. Therefore,

$$V_{\text{eff}} \propto \begin{cases} 1/r^2 & \text{(small } r) \\ -1/r & \text{(large } r) \end{cases} \tag{9.8}$$

If this analysis and sketch look familiar, that's because they are. The effective potential has the same form as the energy in hydrogen (Section 7.3 or `r26-lumping-hydrogen.pdf`); therefore, our conclusions about planetary motion will generalize to hydrogen.

Imagine the planet orbiting in a circular orbit. In that orbit, r remains constant so dr/dt and d^2r/dt^2 are both zero. For that to be true, the particle must live where the effective potential V_{eff} is flat – in other words, at its minimum at $r = r_0$.

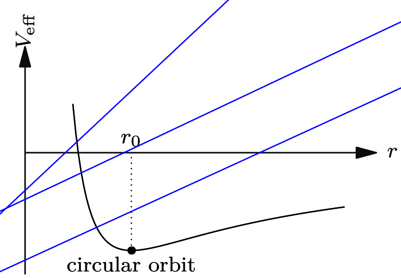

Now perturb the orbit by kicking the planet slightly outwards. That kick does not change the angular momentum ℓ, because angular momentum depends on the tangential velocity. But it gives the planet a nonzero radial velocity ($dr/dt \neq 0$). Thus, it now has r-coordinate kinetic energy (the θ-coordinate kinetic energy is taken care of by the centrifugal-potential piece of the effective potential). The orbital radius r then varies between the extremes where the r kinetic energy turns completely into effective-potential energy. Those are the two points where the horizontal line intersects the effective potential, and the corresponding radii are the minimum and maximum orbital distances.

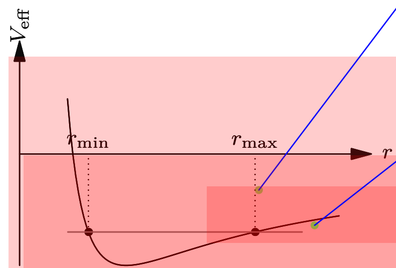

I'm confused as to if this 'kick' is one that knocks the planet into a different type of orbit or if it just shifts the orbit by a certain distance....

I think it just means to disturb it from its original orbit.

Think of a comet

What gives the planet its initial perturbation? That it is nearly impossible to form a planet with 0 initial radial velocity?

does this mean that the orbits were circular until they were disturbed?

Good point.

Can r_min and r_max be anywhere just as long as r_min is less that r0 and r_max is greater than r0?

so are we saying that the dtheta/dt is always constant during this? wasnt sure with he wording

I think this is a little easier to visualize if you say something like, "imaging a marble resting on the line we have drawn and you thump it. It will slide up one side and then up the other. Then it will eventually roll back to the minimum energy position. The farthest points it rolls to form the max and min." This gives a visual and intuitive representation of the conversion from kinetic to potential energy as the ball rolls up the hill.

Thanks for this point, it makes me visualize this situation a lot better.

That's a nice picture. I had talked myself out of using it for the following reason. In the hill picture, the marble's velocity along the hill (not the horizontal component of that velocity) is determined by the kinetic energy (the difference between the total energy and the potential-energy curve). For the planet, the kinetic energy translates directly into a horizontal component (dr/dt). But that flaw may be small compared to the benefit in helping one's intuition.

$1/r^2$ overwhelms the $1/r$ in the gravitational potential. At large r, the gravitational potential is the important term, because its $1/r$ approaches zero more slowly than does the $1/r^2$ in the centrifugal potential. Therefore,

$$V_{\text{eff}} \propto \begin{cases} 1/r^2 & \text{(small } r\text{)} \\ -1/r & \text{(large } r\text{)} \end{cases} \tag{9.8}$$

If this analysis and sketch look familiar, that's because they are. The effective potential has the same form as the energy in hydrogen (Section 7.3 or `r26-lumping-hydrogen.pdf`); therefore, our conclusions about planetary motion will generalize to hydrogen.

Imagine the planet orbiting in a circular orbit. In that orbit, r remains constant so dr/dt and d^2r/dt^2 are both zero. For that to be true, the particle must live where the effective potential V_{eff} is flat – in other words, at its minimum at $r = r_0$.

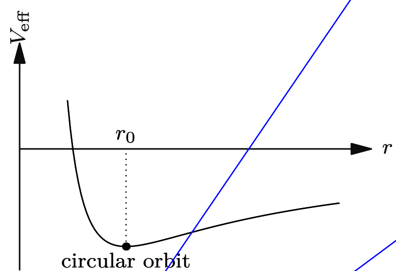

Now perturb the orbit by kicking the planet slightly outwards. That kick does not change the angular momentum ℓ, because angular momentum depends on the tangential velocity. But it gives the planet a nonzero radial velocity ($dr/dt \neq 0$). Thus, it now has r-coordinate kinetic energy (the θ-coordinate kinetic energy is taken care of by the centrifugal-potential piece of the effective potential). The orbital radius r then varies between the extremes where the r kinetic energy turns completely into effective-potential energy. Those are the two points where the horizontal line intersects the effective potential, and the corresponding radii are the minimum and maximum orbital distances.

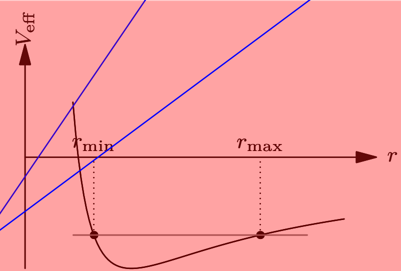

This is a lot of information presented very quickly and for those who have forgotten their 8.01, it's extremely confusing. You might want to expand on this section a bit.

Yeah, agreed. I'm getting a little lost in all the math here that isn't really explained.

Yeah, I am one of those people who has forgotten 8.01 so a longer explanation would definitely be helpful to me.. However, I wonder if a too in-depth would take away from the point of this section... maybe for people who are unfamiliar with this concept there could be an appendix in the back

I'm a bit lost too – where are the springs!?

At first, I expected that after perturbing the orbit, it would return to its lowest energy state, but then I remembered that there is conservation of momentum.

This is really interesting- what type of "kick" could have caused this phenomenon in real life?

I'm not sure you this type of "kick" would ever actually occur.

Maybe collision with a smaller object?

Maybe a large meteor hitting the planet and altering its path a bit? Seems like a kick...

This was the only kick I could think of...would seem difficult to get a significant man-made kick.

What the kick does is to make the angular momentum no longer match with the total energy for a circular orbit. So, although it's hard to imagine the kick directly, one can imagine other ways in which the angular momentum does not match with the total energy for a circular orbit. For example, if the planet is going too fast for a circular orbit, it will move in an ellipse (or even a hyperbola).

can we see this last part in a picture/image of what it means?

But what shape is that orbit? Finding the shape seemingly requires solving the differential equation for r, using the conservation of angular momentum to find $d\theta/dt$, and then integrating to find $\theta(t)$.

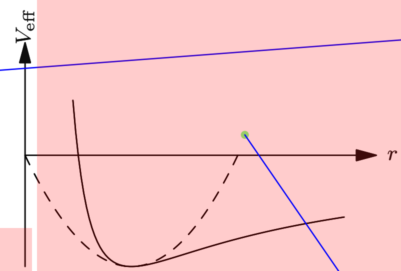

An alternative approach is to use the observation that, near its minimum, the effective potential is shaped like a parabola. We can find that parabola by expanding V_{eff} in a Taylor series about the minimum. To do so, first find the the minimum-energy radius r_0 (the radius of the circular orbit). By setting $dV_{\text{eff}}/dr = 0$, we find

$$r_0 = \frac{\ell^2}{GM}. \tag{9.9}$$

The Taylor series about r_0 is

$$V_{\text{eff}}(r) = V_{\text{eff}}(r_0) + (r - r_0)\left.\frac{dV_{\text{eff}}}{dr}\right|_{r=r_0} + \frac{1}{2}(r-r_0)^2\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} + \cdots. \tag{9.10}$$

The first term in the Taylor series is just a constant, so it has no effect on the motion of the planet (forces depend on differences in energies, so a constant offset has no effect). The second term vanishes because at the minimum energy, i.e. where $r = r_0$, the slope of V_{eff} is zero. The third term contains the interesting physics. To evaluate it, we first need to compute the second derivative:

$$\frac{d^2V_{\text{eff}}}{dr^2} = -2\frac{GM}{r^3} + 3\frac{\ell^2}{r^4}. \tag{9.11}$$

At $r = r_0$, it becomes (after using $r_0 = \ell^2/GM$)

$$\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} = \frac{GM}{r_0^3}. \tag{9.12}$$

Therefore, the Taylor approximation is

$$V_{\text{eff}} = \text{irrelevant constant} + \frac{1}{2}\frac{GM}{r_0^3}(r-r_0)^2. \tag{9.13}$$

Compare this potential energy per mass to the potential energy (per mass) for a mass on a spring,

Comments on page 4

Isn't this why calculus was invented?

But we don't like using calculus in this class!

I'm confused. How are r_min and r_max not just the minor and major axis of the ellipse?

Aren't they though? Though I think Sanjoy is not assuming we know the shape is an ellipse and instead is leaving it open to be any given shape.

it is kind of weird that we are assuming something that we don't already know something

We're not assuming that it's an ellipse. We have, however, deduced that, whatever the shape, it will have an r_min and r_max (set by conservation of energy). Later, once we find out how fast the planet oscillates between r_min and r_max, we'll also see that it's an ellipse with the sun not at the center.

interesting approach to looking at the orbit of a planet..

when should we use each approach? should we just use this one when we want to see the shape?

The effective potential shape is the actual shape of the orbit?

I believe it should be due to the fact that the potential relies on the distance.

The potential here is a graph that tells you the effective potential depending on how far you are from r=0.

it isn't the shape of hte orbit, but it's shape determines the characteristics of the orbit, such as where it is and what its frequency is.

thank you

if only i could remember how to use taylor series...

But what shape is that orbit? Finding the shape seemingly requires solving the differential equation for r, using the conservation of angular momentum to find $d\theta/dt$, and then integrating to find $\theta(t)$.

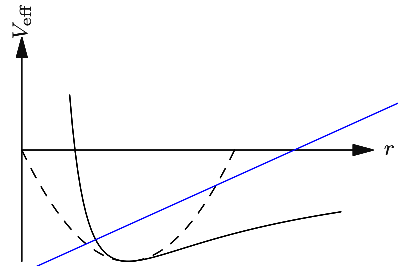

An alternative approach is to use the observation that, near its minimum, the effective potential is shaped like a parabola. We can find that parabola by expanding V_{eff} in a Taylor series about the minimum. To do so, first find the the minimum-energy radius r_0 (the radius of the circular orbit). By setting $dV_{\text{eff}}/dr = 0$, we find

$$r_0 = \frac{\ell^2}{GM}. \tag{9.9}$$

The Taylor series about r_0 is

$$V_{\text{eff}}(r) = V_{\text{eff}}(r_0) + (r - r_0)\left.\frac{dV_{\text{eff}}}{dr}\right|_{r=r_0} + \frac{1}{2}(r-r_0)^2\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} + \cdots. \tag{9.10}$$

The first term in the Taylor series is just a constant, so it has no effect on the motion of the planet (forces depend on differences in energies, so a constant offset has no effect). The second term vanishes because at the minimum energy, i.e. where $r = r_0$, the slope of V_{eff} is zero. The third term contains the interesting physics. To evaluate it, we first need to compute the second derivative:

$$\frac{d^2V_{\text{eff}}}{dr^2} = -2\frac{GM}{r^3} + 3\frac{\ell^2}{r^4}. \tag{9.11}$$

At $r = r_0$, it becomes (after using $r_0 = \ell^2/GM$)

$$\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} = \frac{GM}{r_0^3}. \tag{9.12}$$

Therefore, the Taylor approximation is

$$V_{\text{eff}} = \text{irrelevant constant} + \frac{1}{2}\frac{GM}{r_0^3}(r-r_0)^2. \tag{9.13}$$

Compare this potential energy per mass to the potential energy (per mass) for a mass on a spring,

it would help me if there was a line above this with the equation where this came from

I believe there is, and it's "By setting dV/dr=0..."

All he did was differentiate equation (9.7) with respect to r and set the derivative equal to 0. The value of r that solved that equation he called r_0

is there another method we can use besides Taylor seies? I thought using series is to make things simpler.

I would not/do not remember how to do this. Oops.

Nor do I, but the frequency that Taylor series pop up in my classes probably means I should look at them again...

It may be useful to include an appendix in the book of useful math tools for people like me who forgot about Taylor series...

I agree, a reference of common equations/expansions would be really useful. On a side note, there could also be a related section with all the handy approximation formulas we've derived over the course of the text.

This is a lot to remember. I'm glad we have this written here.

How many terms should we evaluate?

yeah i definitely would not have remembered how to do this :S

For Taylor series and other normally useless math formulas (like Euler), I did what Sanjoy suggested for common variables and wrote them all down on a piece of paper. Not only do I have a reference now, it helps me remember them too.

Having things like Taylor series scares me...

But what shape is that orbit? Finding the shape seemingly requires solving the differential equation for r, using the conservation of angular momentum to find $d\theta/dt$, and then integrating to find $\theta(t)$.

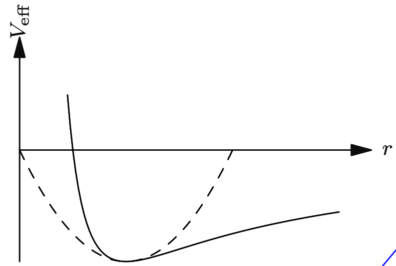

An alternative approach is to use the observation that, near its minimum, the effective potential is shaped like a parabola. We can find that parabola by expanding V_{eff} in a Taylor series about the minimum. To do so, first find the the minimum-energy radius r_0 (the radius of the circular orbit). By setting $dV_{\text{eff}}/dr = 0$, we find

$$r_0 = \frac{\ell^2}{GM}. \tag{9.9}$$

The Taylor series about r_0 is

$$V_{\text{eff}}(r) = V_{\text{eff}}(r_0) + (r - r_0)\left.\frac{dV_{\text{eff}}}{dr}\right|_{r=r_0} + \frac{1}{2}(r - r_0)^2 \left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} + \cdots. \tag{9.10}$$

The first term in the Taylor series is just a constant, so it has no effect on the motion of the planet (forces depend on differences in energies, so a constant offset has no effect). The second term vanishes because at the minimum energy, i.e. where $r = r_0$, the slope of V_{eff} is zero. The third term contains the interesting physics. To evaluate it, we first need to compute the second derivative:

$$\frac{d^2V_{\text{eff}}}{dr^2} = -2\frac{GM}{r^3} + 3\frac{\ell^2}{r^4}. \tag{9.11}$$

At $r = r_0$, it becomes (after using $r_0 = \ell^2/GM$)

$$\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} = \frac{GM}{r_0^3}. \tag{9.12}$$

Therefore, the Taylor approximation is

$$V_{\text{eff}} = \text{irrelevant constant} + \frac{1}{2}\frac{GM}{r_0^3}(r - r_0)^2. \tag{9.13}$$

Compare this potential energy per mass to the potential energy (per mass) for a mass on a spring,

I don't see how the third term is interesting but the 2nd term is not. At minimum energy (r = r0) doesn't the 3rd term and all subsequent terms go to zero as well?

So, all the (r-r0)s are not evaluated at r0, but just the derivatives of Veff. The second term vanishes because the derivative of Veff is 0.

the second derivative is not necessarily 0. if you think about the curve, the first derivative is 0 at the min, but then it starts increasing, so the second derivative is positive.

but the second derivative at r0 is 0, which is where it's being evaluated

if we were just going to use calculus, why do we need to go through all these steps?

Because this is very simple calculus on a problem that was reduced in complexity greatly. I don't think it's uncommon to still use techniques such as calculus, etc. to help us solve problems, but it's about reducing the complexity of these problems so they are easy to solve with simple "second derivatives" as demonstrated here.

what about the other terms in the taylor series? aren't they important? if not why do we ignore them?

As long as the perturbations are small (as long as the orbit is close to a circle), the other terms are even smaller than the quadratic term.

But what shape is that orbit? Finding the shape seemingly requires solving the differential equation for r, using the conservation of angular momentum to find $d\theta/dt$, and then integrating to find $\theta(t)$.

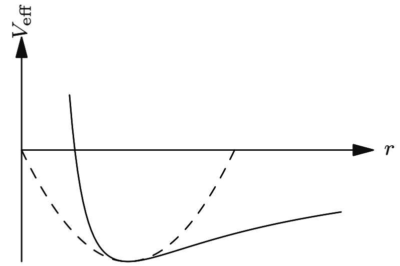

An alternative approach is to use the observation that, near its minimum, the effective potential is shaped like a parabola. We can find that parabola by expanding V_{eff} in a Taylor series about the minimum. To do so, first find the the minimum-energy radius r_0 (the radius of the circular orbit). By setting $dV_{\text{eff}}/dr = 0$, we find

$$r_0 = \frac{\ell^2}{GM}. \tag{9.9}$$

The Taylor series about r_0 is

$$V_{\text{eff}}(r) = V_{\text{eff}}(r_0) + (r - r_0)\left.\frac{dV_{\text{eff}}}{dr}\right|_{r=r_0} + \frac{1}{2}(r - r_0)^2\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} + \cdots. \tag{9.10}$$

The first term in the Taylor series is just a constant, so it has no effect on the motion of the planet (forces depend on differences in energies, so a constant offset has no effect). The second term vanishes because at the minimum energy, i.e. where $r = r_0$, the slope of V_{eff} is zero. The third term contains the interesting physics. To evaluate it, we first need to compute the second derivative:

$$\frac{d^2V_{\text{eff}}}{dr^2} = -2\frac{GM}{r^3} + 3\frac{\ell^2}{r^4}. \tag{9.11}$$

At $r = r_0$, it becomes (after using $r_0 = \ell^2/GM$)

$$\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} = \frac{GM}{r_0^3}. \tag{9.12}$$

Therefore, the Taylor approximation is

$$V_{\text{eff}} = \text{irrelevant constant} + \frac{1}{2}\frac{GM}{r_0^3}(r - r_0)^2. \tag{9.13}$$

Compare this potential energy per mass to the potential energy (per mass) for a mass on a spring,

Just out of curiosity- what is this irrelevant constant?

Maybe the value of the constant is much smaller compared to the value of the second term?

Why do you bother to make that constant irrelevant but then still have the 1/2 in the equation?

The constant is irrelevant because the only thing that matters is the difference between two potential energies, not the absolute value itself (i.e. it doesn't matter what you set to be 0).

This is the case in most types of potential energy we've seen. Take electrical potential for instance – it doesn't matter if you consider ground to be 0, 10, or 100 Volts, as long as the relative voltages are correct.

The constant is irrelevant because the only thing that matters is the difference between two potential energies, not the absolute value itself (i.e. it doesn't matter what you set to be 0).

This is the case in most types of potential energy we've seen. Take electrical potential for instance – it doesn't matter if you consider ground to be 0, 10, or 100 Volts, as long as the relative voltages are correct.

it does seem like a little more explanation is needed here.

Looks kinda similar...

but i dont remember what the potential energy per mass of a mass on a spring looks like...

But what shape is that orbit? Finding the shape seemingly requires solving the differential equation for r, using the conservation of angular momentum to find $d\theta/dt$, and then integrating to find $\theta(t)$.

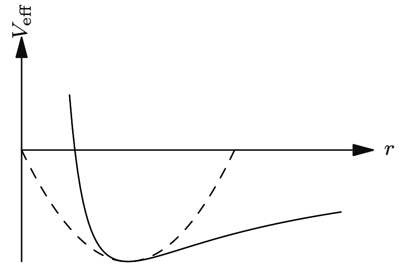

An alternative approach is to use the observation that, near its minimum, the effective potential is shaped like a parabola. We can find that parabola by expanding V_{eff} in a Taylor series about the minimum. To do so, first find the the minimum-energy radius r_0 (the radius of the circular orbit). By setting $dV_{\text{eff}}/dr = 0$, we find

$$r_0 = \frac{\ell^2}{GM}. \tag{9.9}$$

The Taylor series about r_0 is

$$V_{\text{eff}}(r) = V_{\text{eff}}(r_0) + (r - r_0)\left.\frac{dV_{\text{eff}}}{dr}\right|_{r=r_0} + \frac{1}{2}(r-r_0)^2 \left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} + \cdots. \tag{9.10}$$

The first term in the Taylor series is just a constant, so it has no effect on the motion of the planet (forces depend on differences in energies, so a constant offset has no effect). The second term vanishes because at the minimum energy, i.e. where $r = r_0$, the slope of V_{eff} is zero. The third term contains the interesting physics. To evaluate it, we first need to compute the second derivative:

$$\frac{d^2V_{\text{eff}}}{dr^2} = -2\frac{GM}{r^3} + 3\frac{\ell^2}{r^4}. \tag{9.11}$$

At $r = r_0$, it becomes (after using $r_0 = \ell^2/GM$)

$$\left.\frac{d^2V_{\text{eff}}}{dr^2}\right|_{r=r_0} = \frac{GM}{r_0^3}. \tag{9.12}$$

Therefore, the Taylor approximation is

$$V_{\text{eff}} = \text{irrelevant constant} + \frac{1}{2}\frac{GM}{r_0^3}(r - r_0)^2. \tag{9.13}$$

Compare this potential energy per mass to the potential energy (per mass) for a mass on a spring,

I was wondering when we would finally tie this in with springs... maybe you could put a brief outline of how you approach this problem in the beginning of the reading, because so far, I've just been reading this and wondering what this has to do with springs.

Yeah, the title was deceiving.

Really? I just assumed from the moment we started talking about energy we'd have some system acting like a spring whereby PE is converted to KE, and so forth... I think the diagram also sort of hints at that.

I think we can assume those connections, but the lead in is just way too long.

true, but it would still be nice to have a brief outline of the approach before diving in to the very complicated problem

$$V = \frac{1}{2}\frac{k}{m}(x-a)^2, \quad (9.14)$$

where a is the spring's equilibrium length. The spring and the planet have the same form for the potential energy per mass. One needs only the following mapping:

$$x \leftrightarrow r.$$
$$a \leftrightarrow r_0. \quad (9.15)$$
$$k/m \leftrightarrow GM/r_0^3.$$

Therefore, like the x coordinate for the mass on the spring, the planet's distance to the sun (the r coordinate) oscillates in simple harmonic motion! For the mass on a spring, the angular frequency of oscillation is $\omega = \sqrt{k/m}$. Therefore, using the preceding mappings, we find that the planet's radial distance oscillates about r_0 with angular frequency

$$\omega_{\text{perturb}} = \sqrt{\frac{GM}{r_0^3}}. \quad (9.16)$$

The planet's motion is therefore described by two frequencies. The first is ω_{perturb}, the just-computed oscillation frequency of the r coordinate. The second is ω_{orbit}, the oscillation frequency of the orbital motion around the sun (the tangential frequency). Their dimensionless ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$ determines the shape of the orbit. Here are the orbit shapes marked with the ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$, with each orbit drawn against the unperturbed circular orbit.

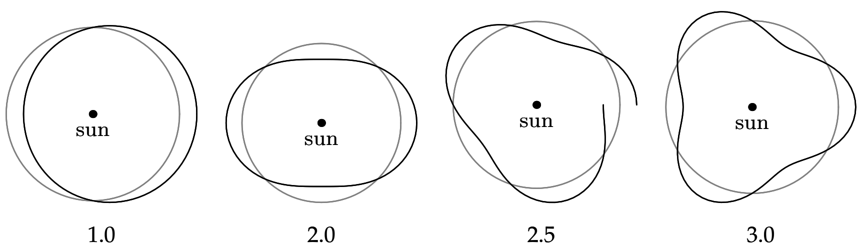

The orbital frequency of the circular orbit is

$$\omega_{\text{orbit}} = \frac{v}{r_0}, \quad (9.17)$$

where v is the orbital velocity (the tangential velocity). To solve for v/r_0, equate the centripetal acceleration to the gravitational acceleration:

would help to put "+ irrelevant constant" here, I think...

Is our familiar formula for the potential energy of a spring an approximation?

We just calculated V_eff in general, and now we want to compare it to potential energy per mass for a spring. They're two different calculations, and this one doesn't have the extra constant.

why is there a comma?

Proper sentence punctuation!

that's so interesting. it makes sense, but i never would have guessed it/believed it/asserted it

The use of the word "one" in this sentence is confusing- I understand that it's trying to say "one" as in a person, but it could also come off as "one" referring to either the spring or the planet and in that case the sentence is extremely confusing.

I often hear the term "one needs only" and immediately understood what he was referring to. In either way I think the sentence conveys the point appropriately.

I like this mapping it makes the analogy really clear

yep I agree, even though I feel like the arrow shouldn't be double sided.

Yeah I kind of stopped on the formula for a little bit trying to think this through only to realize it was explained later. It may be better to have it say something like : If you take these steps and convert the following variables, then you have this equation.

I feel like this is a little overly obvious from the comparison above and doesn't need to be so explicitly stated...

I actually like this. It's nice to have everything all in one place and it's a good summary.

Yeah, I also feel like it's good to have this comparison here. It definitely doesn't hurt, so what motive is there to take it out. I say keep it.

Definitely keep it. I know I'm not alone from the other comments that these couple lines helps a lot in clarifying these relationships.

I love the clarifications in any sort of graphical form!

$$V = \frac{1}{2}\frac{k}{m}(x-a)^2, \tag{9.14}$$

where a is the spring's equilibrium length. The spring and the planet have the same form for the potential energy per mass. One needs only the following mapping:

$$x \leftrightarrow r.$$
$$a \leftrightarrow r_0. \tag{9.15}$$
$$k/m \leftrightarrow GM/r_0^3.$$

Therefore, like the x coordinate for the mass on the spring, the planet's distance to the sun (the r coordinate) oscillates in simple harmonic motion! For the mass on a spring, the angular frequency of oscillation is $\omega = \sqrt{k/m}$. Therefore, using the preceding mappings, we find that the planet's radial distance oscillates about r_0 with angular frequency

$$\omega_{\text{perturb}} = \sqrt{\frac{GM}{r_0^3}}. \tag{9.16}$$

The planet's motion is therefore described by two frequencies. The first is ω_{perturb}, the just-computed oscillation frequency of the r coordinate. The second is ω_{orbit}, the oscillation frequency of the orbital motion around the sun (the tangential frequency). Their dimensionless ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$ determines the shape of the orbit. Here are the orbit shapes marked with the ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$, with each orbit drawn against the unperturbed circular orbit.

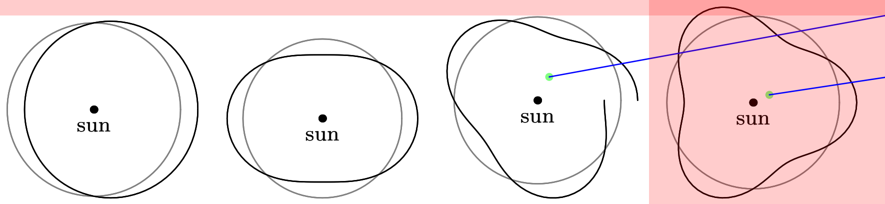

The orbital frequency of the circular orbit is

$$\omega_{\text{orbit}} = \frac{v}{r_0}, \tag{9.17}$$

where v is the orbital velocity (the tangential velocity). To solve for v/r_0, equate the centripetal acceleration to the gravitational acceleration:

This is a good presentation of the material.

this two terms look really different - are they actually comparable?

They do look really different, but when you simplify the dimensions, they are comparable: k/m = (F/d)/m = (ma/d)/m = a/d = 1/t^2. GM/d^3 = (md^3/mt^2)/d^3 = 1/t^2.

That's really cool.

And it ties back in with dimensional analysis! You should mention that here, it seems like a good place to apply our dimensional analysis skills.

you mean if you considered the orbit in 2d or something? a circular orbit would leave the planet at nearly the same distance the entire time...?

This is very interesting - it took a lot of work to get us here though, so I'm still not sure how applicable springs are for generally approximation problems.

I like this connection even though it did take a while to get here. At first I expected we would only be doing approximations that could happen very quickly but after seeing cases like this it makes me understand that you can solve an extremely long problem in a decent amount of time using approximations.

can you explain a little back round about what kind of perturbations are occurring

but I thought that planets orbited in ellipse but not centered around the sun

why isn't this orbit closed?

Are there any planets that actually orbit like this?

$$V = \frac{1}{2}\frac{k}{m}(x - a)^2, \quad (9.14)$$

where a is the spring's equilibrium length. The spring and the planet have the same form for the potential energy per mass. One needs only the following mapping:

$$x \leftrightarrow r.$$
$$a \leftrightarrow r_0. \quad (9.15)$$
$$k/m \leftrightarrow GM/r_0^3.$$

Therefore, like the x coordinate for the mass on the spring, the planet's distance to the sun (the r coordinate) oscillates in simple harmonic motion! For the mass on a spring, the angular frequency of oscillation is $\omega = \sqrt{k/m}$. Therefore, using the preceding mappings, we find that the planet's radial distance oscillates about r_0 with angular frequency

$$\omega_{\text{perturb}} = \sqrt{\frac{GM}{r_0^3}}. \quad (9.16)$$

The planet's motion is therefore described by two frequencies. The first is ω_{perturb}, the just-computed oscillation frequency of the r coordinate. The second is ω_{orbit}, the oscillation frequency of the orbital motion around the sun (the tangential frequency). Their dimensionless ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$ determines the shape of the orbit. Here are the orbit shapes marked with the ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$, with each orbit drawn against the unperturbed circular orbit.

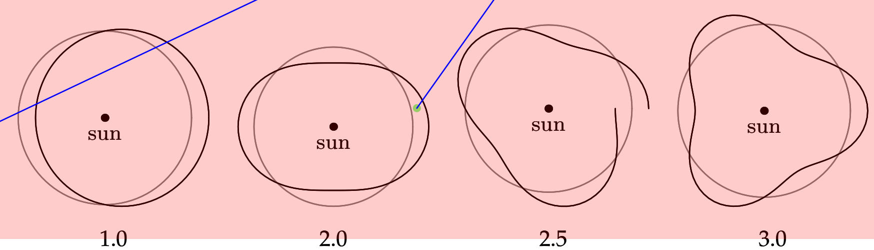

The orbital frequency of the circular orbit is

$$\omega_{\text{orbit}} = \frac{v}{r_0}, \quad (9.17)$$

where v is the orbital velocity (the tangential velocity). To solve for v/r_0, equate the centripetal acceleration to the gravitational acceleration:

This is really cool

While the picture is cool, I think it might be useful to provide a little more background on the ratios...

Agreed, this diagram is awesome. It's really cool to think of an orbit like that. Could you maybe add more to this diagram for higher rations of w_perturb / w_orbit. I think it would be interesting to see this at many different values.

I agree that this is a cool diagram, but as said above, I think it'd be nice to see some more extreme cases of what really high and really low frequencies can look like.

Its interesting how the ratio of frequencies must be a whole number in order for the orbit to be closed

How do you get from the ratio to these shapes? Also, why is 2.5 an open loop? Would a ratio of 4 have 4 lobes?

I agree. I'd like to see more explanation of why these diagrams take the shapes they do.

Yeah, I'm pretty lost about where these pictures came from, and also how you drew these conclusions by applying springs. I think a little more explanation is necessary here.

$$V = \frac{1}{2}\frac{k}{m}(x-a)^2, \quad (9.14)$$

where a is the spring's equilibrium length. The spring and the planet have the same form for the potential energy per mass. One needs only the following mapping:

$$\begin{aligned} x &\leftrightarrow r. \\ a &\leftrightarrow r_0. \\ k/m &\leftrightarrow GM/r_0^3. \end{aligned} \quad (9.15)$$

Therefore, like the x coordinate for the mass on the spring, the planet's distance to the sun (the r coordinate) oscillates in simple harmonic motion! For the mass on a spring, the angular frequency of oscillation is $\omega = \sqrt{k/m}$. Therefore, using the preceding mappings, we find that the planet's radial distance oscillates about r_0 with angular frequency

$$\omega_{\text{perturb}} = \sqrt{\frac{GM}{r_0^3}}. \quad (9.16)$$

The planet's motion is therefore described by two frequencies. The first is ω_{perturb}, the just-computed oscillation frequency of the r coordinate. The second is ω_{orbit}, the oscillation frequency of the orbital motion around the sun (the tangential frequency). Their dimensionless ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$ determines the shape of the orbit. Here are the orbit shapes marked with the ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$, with each orbit drawn against the unperturbed circular orbit.

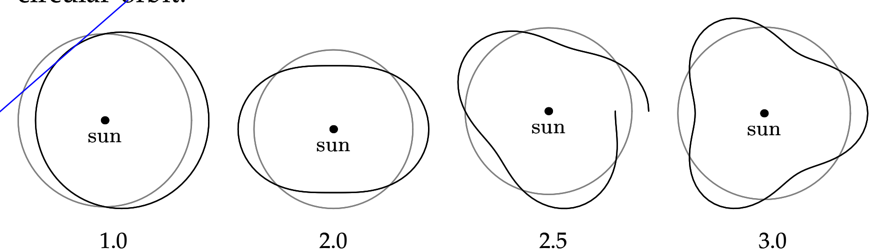

The orbital frequency of the circular orbit is

$$\omega_{\text{orbit}} = \frac{v}{r_0}, \quad (9.17)$$

where v is the orbital velocity (the tangential velocity). To solve for v/r_0, equate the centripetal acceleration to the gravitational acceleration:

How do you get from the ratio to these shapes? Also, why is 2.5 an open loop? Would a ratio of 4 have 4 lobes?

Yeah, it might be helpful to see a ratio of 4 if it's possible to fit it within this space... But I do think the ratio is directly correlated to the number of lobes.

I'm confused too. These need more explanation, if they're really necessary.

yeah i'm also curious about 2.5. after a certain number of orbits does the path eventually start repeating itself?

I believe 2.5 is open because it is not a whole number. After a whole number, a rotation is finished, and it can line up again to complete the circuit. However, after only half a rotation, the the orbit misaligns.

you get from the ratio to these shapes because the ratios are how many periods are there in the period of a normal, circiular orbit.

While I don't feel like these kinds of shapes exist in real life, it is essentially saying this:

Take your normal circular orbit. One full circle is one period. If we stretch that length out in a straight line, that is a smooth, undisturbed movement.

Then, we add a disturbance in the form of a sinusoidal oscillation (this is modeled after springs after all). In the case that the period of the disturbance is exactly the same as the period of the circular orbit, half the sinusoidal disturbance will deviate downwards from the [straight line, that represents the stretched out circular orbit] and half will be upwards. If we overlap this new sinusoidal path on top of the straight line of the straightened out circular path, and then reconnect the straight line path back into a circle, we will get the first shape.

(This is kind of like the concept of making mobius strips out of paper. It would've been more intuitive, in my opinion, if he stretched out the paths first and then rejoined them)

Another way of looking at this problem is to think of the sinusoidal disturbance as a *radial* force that varies sinusoidally. This means that at any given point along the circle, the disturbance can either point inwards towards the center, or outward parallel to the radius. In this way, you can imagine that only disturbance frequencies that are whole number multiples of the frequency of the original circular orbit will line up where it started. As a simple check: In the first shape, the disturbance frequency is exactly the same (multiple of 1) as the circular orbit's frequency. Therefore, a sinusoidal disturbance will spend the first half of the orbit pushing inwards, and the second half pushing radially outwards. The net effect is to shift the entire circle to the right.

$$V = \frac{1}{2}\frac{k}{m}(x-a)^2, \quad (9.14)$$

where a is the spring's equilibrium length. The spring and the planet have the same form for the potential energy per mass. One needs only the following mapping:

$$\begin{aligned} x &\leftrightarrow r. \\ a &\leftrightarrow r_0. \\ k/m &\leftrightarrow GM/r_0^3. \end{aligned} \quad (9.15)$$

Therefore, like the x coordinate for the mass on the spring, the planet's distance to the sun (the r coordinate) oscillates in simple harmonic motion! For the mass on a spring, the angular frequency of oscillation is $\omega = \sqrt{k/m}$. Therefore, using the preceding mappings, we find that the planet's radial distance oscillates about r_0 with angular frequency

$$\omega_{\text{perturb}} = \sqrt{\frac{GM}{r_0^3}}. \quad (9.16)$$

The planet's motion is therefore described by two frequencies. The first is ω_{perturb}, the just-computed oscillation frequency of the r coordinate. The second is ω_{orbit}, the oscillation frequency of the orbital motion around the sun (the tangential frequency). Their dimensionless ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$ determines the shape of the orbit. Here are the orbit shapes marked with the ratio $\omega_{\text{perturb}}/\omega_{\text{orbit}}$, with each orbit drawn against the unperturbed circular orbit.

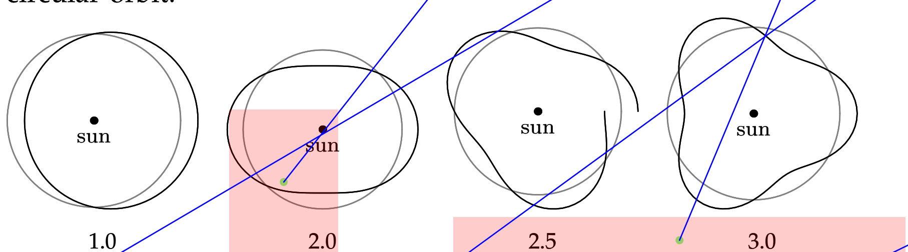

The orbital frequency of the circular orbit is

$$\omega_{\text{orbit}} = \frac{v}{r_0}, \quad (9.17)$$

where v is the orbital velocity (the tangential velocity). To solve for v/r_0, equate the centripetal acceleration to the gravitational acceleration:

The diagram at the end looks more like a circle while I would consider this to resemble an ellipse better. How does the bigger w(perturb) make the orbit shape more ellipse-y?

are there actually any orbits that look like this?

If you're asking has this happened, then yes. If you're asking is there a body that has a stable orbit like this, no, because these are perturbances in the orbit, not forces constantly acting on them during every period.

Out of curiosity, when has it happened before?

I'm sure tons of space debris has been subject to high perturbation to orbit ratio. However I doubt they stay in it very long.

so is this just turbulence in the orbit that happens occasionally until it stabilizes?

If you keep reading, it says that the ratio is one, so it would seem that the answer is 'no'.

It would be nicer to have the definition of omega(orbit) before the diagrams where you are using the ratio

somewhat related question: in elliptical orbits the velocity is greater at certain points around the ellipse...is this velocity that much greater and at what point does it happen for Earth?

It remember it has to do with sweeping out equal areas in equal time. So when it is farther away from the sun, the velocity is slower and when it is closer it is faster. I don't have quantitative numbers though.

That's right. The velocity is only much greater if the orbit is highly elliptical...so for a comet it might make a difference, but not for a planet.

thanks!

this probably should have brought up a lot earlier.

$$\frac{v^2}{r_0} = \frac{GM}{r_0^2}. \quad (9.18)$$

To manufacture v/r_0 on the left side, divide both sides by r_0 and take the square root:

$$\underbrace{\frac{v}{r_0}}_{\omega_{\text{orbit}}} = \sqrt{\frac{GM}{r_0^3}}. \quad (9.19)$$

This expression is also ω_{perturb}. Therefore, $\omega_{\text{orbit}} = \omega_{\text{perturb}}$.

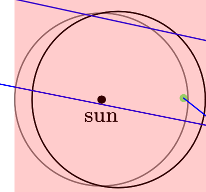

This result explains the elliptical (Kepler) orbits. First, the ratio is an integer, so the orbit is closed (compare the orbit with the ratio 2.5) – as it should be. Second, the ratio is 1, so the orbit's center is slightly away from the sun – as it should for a planetary orbit (the sun is at one focus of the ellipse, not at the center). The surprising conclusion of all the analysis is that a planetary orbit contains a spring; once this fact is appreciated, a spring analysis allows us to understand the complicated orbital motion without solving complicated differential equations.

As a bonus, the effective potential has the same form as the energy in hydrogen. Therefore, that energy also looks like a parabola near its minimum. The consequence is that a chemical bond acts like a spring (for small extensions). This second application is not a mere coincidence. Near a minimum, almost every function looks a parabola, so almost every physical system contains a spring. Springs are everywhere!

is this the right term? it makes it sound like you're making things up.

I agree this sounds funny. But I think the message get through that you're trying to get V/r0 on the left.

yeah I feel like a better word might be "obtain"

The distortion of space-time causes the orbital and angular frequencies to be different, so the orbit of Mercury doesn't perfectly close on itself in a noticeable way.

This is also exactly why the period for a person falling through a frictionless tunnel all the way through the earth is the same as the orbital period of someone just above the surface of the earth, which is a rather random fact we proved in 8.012.

I like this derivation and the explanation that follows. It's pretty cool.

I like you put the visual representation first so the reader gets an idea where the proof is going

I think it would be really cool if you had an experiment of physical phenomena that shows how this orbit is a spring

Either an experiment or some more concrete example that we could reference in the future. It seems like those really help to drive home the points made in class

I'm hoping to see those in the rest of this chapter!

This looks like a circle, not an ellipse to me. Did I miss something?

whether this is a circle or an ellipse, I don't know for sure.

But, circles are a subset of ellipses (they are the special case where the two foci of the ellipse lines up). So saying that this is an ellipse cannot be wrong. But whether it is a circle, i'm not sure.

I think that here it is drawn as a circle but it wouldn't be a circle in reality because the gravity would be a little weaker as the earth gets further from the sun so the orbit would be more of an elongated ellipse. Maybe the circle was just used because it's a simplification that shows where the offset of the center of the orbit.

Oh, I see. Good explanation.

$$\frac{v^2}{r_0} = \frac{GM}{r_0^2}. \tag{9.18}$$

To manufacture v/r_0 on the left side, divide both sides by r_0 and take the square root:

$$\underbrace{\frac{v}{r_0}}_{\omega_{\text{orbit}}} = \sqrt{\frac{GM}{r_0^3}}. \tag{9.19}$$

This expression is also ω_{perturb}. Therefore, $\omega_{\text{orbit}} = \omega_{\text{perturb}}$.

This result explains the elliptical (Kepler) orbits. First, the ratio is an integer, so the orbit is closed (compare the orbit with the ratio 2.5) – as it should be. Second, the ratio is 1, so the orbit's center is slightly away from the sun – as it should for a planetary orbit (the sun is at one focus of the ellipse, not at the center). The surprising conclusion of all the analysis is that a planetary orbit contains a spring; once this fact is appreciated, a spring analysis allows us to understand the complicated orbital motion without solving complicated differential equations.

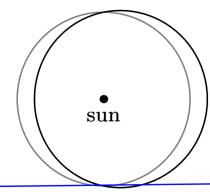

As a bonus, the effective potential has the same form as the energy in hydrogen. Therefore, that energy also looks like a parabola near its minimum. The consequence is that a chemical bond acts like a spring (for small extensions). This second application is not a mere coincidence. Near a minimum, almost every function looks a parabola, so almost every physical system contains a spring. Springs are everywhere!

I don't see why this is necessarily true...

The sun is located at one of the focii of the ellipse that describes the orbit of a planet. The focus of an ellipse is not in the geometric center but on the major axis of the ellipse.

Yes, but why does it make sense that it is this way rather than being a circular orbit with the sun at the center?

The focus of the sun-earth ellipse is very close to the center of the sun because it's so much more massive than the earth (The focus is within the radius of the sun if I recall correctly). This is why we generally think of us orbiting around the sun, but actually its around this focii point.

I understand all the parts individually, but I don't completely understand it when we put it all together... Like, I know that the orbit of the earth is elliptical, but I don't understand quite how we got that conclusion... The diagrams helpa bit but I don't understand where the diagrams came from.

is there a term for the other focus of the ellipse then? does this hold any significance?

stll dont fully get the spring thing- I think this is hard to read and understand for me

after this whole thing i still dont really understand the connection. maybe an example using just springs and not the complicated stuff would help to explain. i suspect that's coming in future sections though.

While I think that the lead in was a bit heavy, the diagrams and simple matching techniques made it clear how a spring and its properties could be applied here.

well it wasnt the easiest set of equations to get to that point either. at least to people unaccustomed to planetary motion

it's so cool how these things are all interconnected

That is awesome.

Agreed.

It's really interesting that planetary orbits, hydrogen, and mass on spring all relate so well. I guess this goes to show that a few physical laws can go very far in explaining the universe

$$\frac{v^2}{r_0} = \frac{GM}{r_0^2}. \tag{9.18}$$

To manufacture v/r_0 on the left side, divide both sides by r_0 and take the square root:

$$\underbrace{\frac{v}{r_0}}_{\omega_{\text{orbit}}} = \sqrt{\frac{GM}{r_0^3}}. \tag{9.19}$$

This expression is also ω_{perturb}. Therefore, $\omega_{\text{orbit}} = \omega_{\text{perturb}}$.

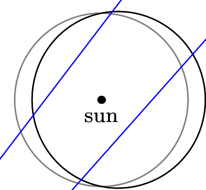

This result explains the elliptical (Kepler) orbits. First, the ratio is an integer, so the orbit is closed (compare the orbit with the ratio 2.5) – as it should be. Second, the ratio is 1, so the orbit's center is slightly away from the sun – as it should for a planetary orbit (the sun is at one focus of the ellipse, not at the center). The surprising conclusion of all the analysis is that a planetary orbit contains a spring; once this fact is appreciated, a spring analysis allows us to understand the complicated orbital motion without solving complicated differential equations.

As a bonus, the effective potential has the same form as the energy in hydrogen. Therefore, that energy also looks like a parabola near its minimum. The consequence is that a chemical bond acts like a spring (for small extensions). This second application is not a mere coincidence. Near a minimum, almost every function looks a parabola, so almost every physical system contains a spring. Springs are everywhere!

I have heard this before. I guess I never really thought about this before, but now I can see better why.

Yeah, I love this analogy. It is really helpful in understanding small, medium, and large scale processes!

This is interesting, but I'm not seeing the benefits of using the spring analysis instead of lumping for the hydrogen problem.

Well, we did end up using the lumping analysis instead of a using the SHO for hydrogen.

I agree that this is interesting. I'm also interested to see just how far this statement holds.

Is there a relevant example of a minimum where modeling as a spring fails horribly?

I suspect that as long as your effective potential is roughly parabolic near it's minimum, then it's fair game to approximate as a spring, and for most potentials this is probably fair near the minimum, but if you had some strange potential that wasn't smooth at it's minimum or something the model might break down.

Yeah I'd say as long as your function is smooth at the minimum you should be able to approximate it this way

I like the analysis that I read somewhere about how basically all atomic bonds can be considered as springs, and heat is just those bonds flexing. I think deffinitly this statement explains a lot.

This was an interesting lecture and pretty easy to follow.

I agree that this was a fairly easy reading to follow, but it wasn't as good as the section on probability. I guess that's my personal preference though, which is neither here nor there.

I totally disagree. it might be that i've done little-to-no work on any of this kind of stuff since 8.01 in 2004, but i found this section quite hard to follow. the outline was good and easy to follow, but I would have liked a lot more on were a lot of this stuff came from.

Interesting (though presented in most basic physics classes, I believe) and simple. I wonder where we're going from here with the spring method.

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

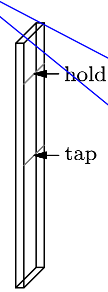

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

I feel like you covered a variety of examples and this is the first time involving music. I am quite happy.

For Wednesday (memo due at 9am on Wed), read about the wood-block demonstration from lecture. Then we'll talk more about physics and music (provided I'm not in the delivery room at Mt Auburn hospital).

I like this introduction. It seems abstract, but you know it is building up to a meaningful conclusion

This was a cool example to see in class! I really like how the readings tie to the lectures, but I also think even for people who don't take 6.055 (which will probably be the case for most readers of this text), that the examples are so interesting to read about and the explanations are very clear.

I know, I'm ashamed I couldn't remember the results of this from streetfighting math..

I like the idea of using springs as music, but the more classical idea of springs. I wonder how that would work...not just vibrational motion but like a slinky

I see that this could be a good way to make the point, but do you really expect anyone to have a carpenter or lumber yard worker make this for them?

I realy don't think that's relevant...

Haha yeah pretty irrelevant

Also, it's not really construction. Most places will happily cut down lumber for you.

Rather than instructing the reader in how to acquire the wood blocks you could simply ask them to acquire blocks - if you're target audience is MIT students I would assume they are capable of figuring out a way to get two wood blocks.

I also assume that this statement applies to the average person as well!

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

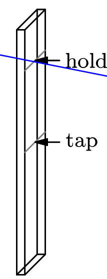

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

I think it would be helpful to say that they are long thin planks rather than blocks, which i would say is a more accurate description of the geometry

I agree...one of the key points from the experiment done in class was that the blocks were identical except for a slight difference in the thickness.

This is interesting, but I have no idea where this is going. This may be a good thing though if explained well, since musical examples are usually lost on me

if they're the same wood plank, wouldn't they have the same thickness? this is a misleading statement.

If this were instead dealing with metal, would the oppose hold true? would the larger piece have a lower frequency?

Maybe instead of including the exact dimensions here, add them into the diagram?

i think you can skip the stuff before this, with the details of asking someone to construct a block for you, and their dimensions. this sentence is sufficient and brings the focus the difference between the two blocks, the real item in question here

I agree. I find the "ask a friend" thing to be a little silly. To me, that's like saying on a homework problem, "first go to the store and buy some paper and a pencil to write down your answers with..."

This is being presented as a home experiment. What kind of Mickey Mouse experimenter would skip details about dimensions and construction? This sentence is just a brief expansion explanation of the above numbers.

I agree with all of this; you might just want to start with: "I have two wood blocks of dimensions...".

I like starting with examples though, it gives me something to picture as I'm going through the reading and it makes it more interesting.

I agree one of the things i like about this book is it's not all fact fact fact. There's actual transition and side stories/information that make it much more interesting to read.

Well, I actually thought it was kind of amusing, but I can see how it is detracting.

It's only 3 sentences...I feel like a 3 sentence intro is completely fine.

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

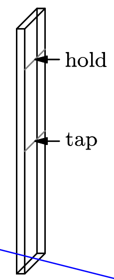

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

How much of a difference does the strength at which you hold it end up mattering? Shouldn't it be the same, as long as you do it the same for both blocks?

I would think yes. I think what he's trying to say is that you need to hold the blocks lightly enough to allow the blocks to vibrate as freely as possible. If you hold them with enough pressure, you could dampen the vibration enough as to make the experiment worthless.

Yeah I'm pretty sure that's what he's saying as well - you don't want to add damping to the vibrations.

I hold them at the "node" – the spot that doesn't move, so it is less sensitive to my pressure.

Maybe you could minimize the human error from holding it by attaching the block to a string and holding the string in your hand, maybe you would have to drill a hole in the wood to thred the string through.

As a phrasing issue, I would suggest that "lightly" be replaced by "loosely."

Maybe draw a hand here to show where you hold it - "toward the edge" seems ambiguous to me.

I agree, especially since holding the block at the "node" will decrease the sensitivity to the pressure of holding it.

How far toward the edge though? How can we tell where the "node" is?

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

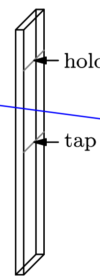

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

When asked to think about this casually, I would have said lower. When asked to think about it as an engineer, I would have said higher. Is it weird that I still separate those mental states?

I did the same thing. There seem to be competing preconceptions here - like bigger objects tend to make deeper sounds but thinner objects can vibrate more easily. Even though I followed the logic in the reading, if you were to ask me a related problem in the future my instincts would still be divided.

I originally thought lower, but thinking of guitar luthiery, one constructs a thicker top to get more clarity in the tone, which is typically understood to be more of the higher overtones and fewer low overtones, a direct correlation to this question. Also interesting is the use of a "tap tone" in guitar construction to correctly voice a top.

Don't know if my train of thought is right, but my first reaction was that the thicker would create a higher pitch, since it vibrates faster, but then I thought about playing a violin and how the highest pitch string is the thinnest. But then I thought more, and the thinnest one is strung tightest, so vibrates fastest still...

If you don't refer to "fundamental frequency" later in the text, is this parenthetical necessary?

The fundamental frequency is a very specific quality of a subject. It may not be entirely necessary, but it is appropriate and accurate.

Is this because what we hear is the fundamental frequency, and not overtones?

The period of sound with a fundamental frequency is the smallest repeating unit of a signal. Overtones, on the other hand, are resonant frequencies that are higher than the fundamental frequency.

You hear the overtones as a "coloring" of the fundamental. That's why a harpsichord and a piano playing the "same" note (the same fundamental) sound like the note is the same pitch but sounds different (the harpsichord sounds more metallic, because it's overtones are stronger).

Actually, the fundamental frequency is so powerful that you hear the fundamental even if it is not there (that's how telephone receivers, or small portable radios, sound at all okay).

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

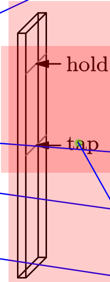

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

I really like this graphic. For some reason, I found the verbal description of this a bit confusing, but the picture really clears up what you mean.

I agree, if you didn't see the demonstration, it may be unclear where the "center" of the board is exactly.

Agreed. And I didn't even see this figure until after I read the entire page. It does reduce some confusion.

The graphic does help a lot, would it be possible to insert some hand shaped image where you ask the user to hold the block?

It's going to be lower, right

I think you missed lecture that day :-P We already did the experiment in class. My reaction was to expect lower too, but the thicker piece was actually higher pitch.

Yeah, I was surprised by that too.

My reaction was that a rigid, stiffer object will vibrate at a higher frequency (pitch). A thicker piece of wood will be stiffer than a thin piece of wood of the same length.

i would say lower

does this have anything to do with the moment of inertia?

I would have been interested to see more examples relating to music this term.

I completely agree. I hadn't really thought about sounds/music as examples for approximation, but in some cases it may be more intuitive than some of the drag problems.

Can you explain "fundamental frequency"? Is it a important term to understanding this?

This is a somewhat weird angle for viewing the part.

I agree...while all the diagrams we have seen in past sections greatly contribute to having a better mental image, this one doesn't seem to bring as much to the table. If anything, using a diagram might trick the reader into thinking this experiment is more complex than it really is.

i think that a diagram is still useful, but it's hard to see how this is a plank. i also think there's not very much information being conveyed by 'hold' and 'tap.' perhaps the distances or other relevant dimensions should be marked.

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

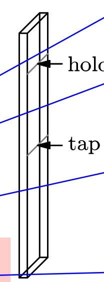

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

Also, will they actually be related? As in a third or fifth from each other, or is it random despite the almost-identicle dimensions?

Did you do this for us in 2.003? I kind of remember it with a metal ruler but I might be making it up

should be different right? kind of related to the stress factor of the board....like in tkd!

I thought it was cool in class when you said that some student could tell which block had the higher frequency just by hearing each block placed on a surface!

This brings up a point I hadn't considered before. How legitimate are these estimation predictions in proposing a scientific hypothesis?

I think the thicker one will produce a lower tone - no real solid reason why, but it feels right- you know?

I think that being able to predict is much better than explaining the result. Predicting demonstrates a true understanding of the physics, whereas any result can be backtracked and justified. Understanding the results is better than nothing though.

I think that being able to predict is much better than explaining the result. Predicting demonstrates a true understanding of the physics, whereas any result can be backtracked and justified. Understanding the results is better than nothing though.

In this case, a random prediction has a 1/2 chance of being correct, whereas understanding the results yields a much higher chance of know what caused them.

I think either method works fine. The main idea is that we apply to our understanding of the world to our observations.

I think also it would be nice if we could think about a slightly different scenario of this problem like knocking on a very thick door and a thin one and comparing the sounds.

I was going to comment on the wording in this sentence; shouldn't there be a 'however' or similar in the beginning of this sentence?

I think there should be...

resonant object could be defined.

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

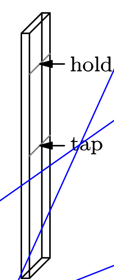

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

I guess because I immediately assumed thicker block = more stiffness, and guessed correctly, it's hard for me to understand some of the counter argument. In particular, what is the logic that a thicker block would mean a longer wavelength?

this is unnecessary considering the picture doesn't explain anything.

The picture maybe should show the 2 blocks, and why the wavelengths would change depending on thickness. (or a diagram to show the effects of the different proposed models, this is something I always have a hard time visualizing)

why isn't this part right? This is what came to mind in class.

It seems to be that the increased stiffness outweighs the other differences.

what if the other two dimensions are also not the same? how would you then compare the wavelength? does the wavelength only depend on the thickness?

Ah, in class today we're going to look at a xylophone and answer that very question.

this is what I originally thought. maybe the reason that the thicker block has a higher not is because it's easier to travel trough the block than through the air so the frequency is higher from the thick wood?

I thought the opposite but when I tried to reason through it I couldn't come up wit ha great answer. I'm assuming in a page or two I'll have my reason

yeah this is how I originally thought about the problem when we were doing the experiment in class.

How can you determine that the thick block is stiffer? I think it might be helpful to add a equation or explain the context of stiffness in greater detail.

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

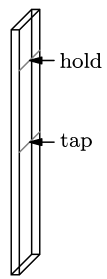

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

i often think to myself something along these lines and confuse myself. i specifically remember this happening a lot on the diagnostic.

I'm not an expert in material science but is this an appropriate model to use for wood?

I thought of the blocks by the first method mentioned. What is a good way to train intuition to use the spring/stiffness method above?

I intuitively used the spring model based on experience with metals of various thickness, however I think the point here is that it's really hard to tell what the right answer is here just using intuition, without doing any calculations.

Just like everyone who commented before me, I have the wrong intuition for thinking about this. I thought that the thicker wood has a longer wavelength, so lower pitch. I need to fix my intuition

I think my source of intuition for this problem was using a xylophone. On a xylophone, the longer bars correspond to low notes, so I just figured more wood = lower pitch.

(addendum: but I guess thickness and size don't exactly correspond well? hm...)

Yea so I had the same wavelengths thoughts, and I also thought about knocking on a wood door and I remember that sound as being lower than the first thin block we heard, but I guess maybe a door is too big to compare to the blocks.

Don't feel bad! When I ask this question in talks I give at physics departments, most of the faculty expect the thicker block to have the lower pitch (because thicker means longer wavelength).

This didn't cross my mind at all. I feel like when it comes to stiffness, the density/mass of the wood overwhelms the effects on the vibrations. My feeling is the stiffness difference is very minute compared to the difference in mass.

I don't understand how stiffness is related to vibration speed. I remember hearing this in class, but we never went over the reasoning.

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

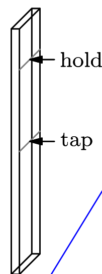

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

I really like that you give a reasoning for each possible argument, it allows each reader to hear a reasonable answer and make their hypothesis before you tell us what prediction is right. It makes even a book seem interactive!

> yeah I agree...it was really helpful because the first argument was exactly what I had thought about in class on Monday...then it was interesting to read about the second argument as a reason for the frequencies to be the same, since I had originally eliminated this possibility without thinking about it very much.

When we say it vibrates faster, does that imply that the amplitude of this motion is smaller? Because when I imagine a thicker, stiffer block, I would assume that it wouldn't move as much in starting amplitude as would a thinner block, and since it says here that the thicker block vibrates faster, I thought that this implied that its starting amplitude was consequently smaller.

> There's a difference between frequency and amplitude. Here we are talking about frequency. sin(theta) and 2*sin(theta) have different amplitudes but the same frequency for example.
>
> > you're right the thicker block will be harder to move, which means it will try to return to it's initial, lowest energy, position faster.
> >
> > > right, so that's an indication of frequency, and not amplitude. amplitude is only how loud the sound is, not what pitch it is at.
> > >
> > > > To be fair, it should in fact have lower amplitude (in addition to higher frequency), assuming you instill the same amount of energy when you tap it.
> > > >
> > > > (since energy increases with both frequency and amplitude, if we increase one we need to decrease the other to keep the energy the same)

Since we are comparing mass and stiffness..please add a equation

I do not think this is the case just looking at things like a xylophone where the slats have slightly different lengths and they change frequency so I think that they would be different

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

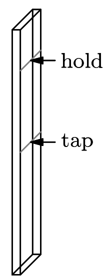

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

I would be surprised if everything canceled. That was the first answer choice that I eliminated in class.

So many other factors have canceled in previous examples that I think it is fair to include that here.

I feel like this should be reworded to something along the lines of "leaving the frequency of the two blocks the same".

It always surprises me how frequently these extra factors seem to cancel themselves out, when it seems so unlikely. However I think it's still a valid hypothesis worth suggesting, even if it doesn't seem very probable.

I like how this paragraph is almost determined to confuse the reader. It gives the analysis more credit knowing that the reader is unsure of what to expect.

This is the exact mental thought process i went though, but until you did the demo in class (and since i already knew the answer) i already knew which term wins out...but still cool!

damn! I was wrong

I would have guessed the thickness of the block outweighed the stiffness, making it a lower pitch

I agree. Prior to the experiment in class, my initial thought was that it was lower. Very much counter-intuitive.

When I heard the problem, I first thought of bending beams and their stiffness. Perhaps because I've taken 2.001 and 2.671. Is it possible that these blocks have harmonics from other bending modes?

I'm not sure if you've taken 2.002, but that class actually derives the physics of this situation. There was one lab where we actually used the frequency of a cantilever beam oscillations to calculate the pitch that would be generated when striking the object.

does the validity of the models depend on the delta magnitude of the thickness?

which model was this? i'm confused by the terminology.

I agree. Resonant-cavity??

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. $30\,\text{cm} \times 5\,\text{cm} \times 1\,\text{cm}$; and
2. $30\,\text{cm} \times 5\,\text{cm} \times 2\,\text{cm}$.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

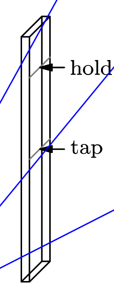

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

Should the fact that we're calling it a resonant _cavity_ be a clue that it is an incorrect model, as we are dealing with the presence of matter (the wood block) rather than the absence (a cavity)?

I think the reason why he keeps calling it a resonant-cavity is that resonant cavities can be modelled using any physical resonator such as an RLC circuit. So maybe the term just refers to a general resonator? maybe?

yeah i was a little confused by this wording too...

Would not have thought this

Me neither, but I still wasn't sure enough to have thought the other answer was correct.

I don't think I would know where to start when trying to come up with this on my own.

I definitely didnt know where to start on this one, but my gut told me to guess higher. I'm excited to see the explanation for this

This is a little confusing because I don't know how to connect the models you proposed your results (most likely due to terminology)

this is an awkward sentence. we did the experiment so "probably" seems a little strange here. we know it's higher for some reason, and the stiffness was our logical argument.

Probably? If the results of the experiment are as you say, why "probably?"

We haven't come up with any mathematical justification for it yet. Thus, we can't assume that because A is false, B is true. The experiment must be examined.

False can imply everything, as we know. If 1 = 2, the I'm the pope.

Are we expected to know that springs model this situation well? You just sort of declare that springs work well in this situation, but I would have no idea if that is true or not.

I think he is first making the statement and then explains it afterward

I agree that this transition is sudden. It seems you were discussing two models in the preceding paragraph and it feels like you've discarded them.

9.2 Musical tones

9.2.1 Wood blocks

Here is a home musical experiment that illustrates proportional reasoning and springs. First construct, or ask a carpenter or a local lumber yard to construct, two wood blocks made from the same larger wood plank. Mine have these dimensions:

1. 30 cm × 5 cm × 1 cm; and
2. 30 cm × 5 cm × 2 cm.

The blocks are identical except in their thickness: 2 cm vs 1 cm.

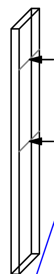

Now tap the thin block at the center while holding it lightly toward the edge, and listen to the musical note. If you do the same experiment to the thick block, will the pitch (fundamental frequency) be higher than, the same as, or lower than the pitch when you tapped the thin block?

You can answer this question in many ways. The first is to do the experiment. It would be nice either to predict the result before doing the experiment or to explain and understand the result after doing the experiment.

One argument is that the block is a resonant object, and the wavelength of the sound depends on the thickness of the block. In that picture, the thick block should have the longer wavelength and therefore the lower frequency. A counterargument, based on a different model of how the sound is made, is that the thick block is stiffer, so it vibrates faster. On the other hand, the thick block is more massive, so it vibrates more slowly. Perhaps these two effects – greater stiffness but greater mass – cancel each other, leaving the frequency unchanged?

I'll do the experiment right now and tell you the result. The thick block has a higher pitch. So the resonant-cavity model is probably wrong. Instead, the stiffness probably more than overcomes the mass.

A spring model explains this result and even predicts the frequency ratio. In the spring model, a wood block is made of wood atoms connected by

Awesome phrase. You might want to call attention to it like you did in class just to point out that you're aware that wood is not an element..

Haha, I like this idea of wood atoms.

I was confused by this in class - how would this model work? how do you know how many "atoms" contribute to the spring ?

it's not a matter of how many, it's just the fact that the atoms are constantly moving creating a spring-like motion and response.

chemical bonds, which are springs. As the block vibrates, it takes these shapes (shown in a side view):

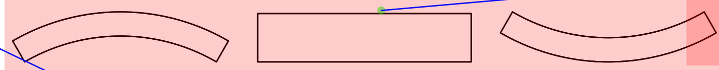

The block is made of springs, and it acts like a big spring. The middle position is the equilibrium position, when the block has zero potential energy and maximum kinetic energy. The potential energy is stored in stretching and compressing the bonds. Imagine deforming the block into a shape like the first shape:

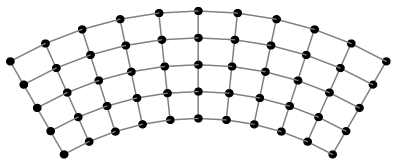

Each dot is a wood 'atom', and each gray line is a spring that models the chemical bond between wood atoms. Deforming the block stretches and compresses these bonds. These numerous individual springs combine to make the block behave like a large spring. Because the block is a big spring, the energy required to to produce a vertical deflection is proportional to the square of the deflection:

$$E \sim ky^2, \tag{9.20}$$

where y is the deflection, and k is the stiffness of the block.

Intuitively, the thicker the block, the stiffer it is (higher k). The spring model will help us find how k depends on the thickness h. To do so, imagine deflecting the thin and thick blocks by the same distance y, then compare their stored energies E_{thin} and E_{thick} by forming their ratio

$$\frac{E_{\text{thick}}}{E_{\text{thin}}}. \tag{9.21}$$

That ratio is

$$\frac{k_{\text{thick}}y^2}{k_{\text{thin}}y^2} = \frac{k_{\text{thick}}}{k_{\text{thin}}} \tag{9.22}$$

because y is the same for the thick and thin blocks. So, the ratio of stored energies is also the ratio of stiffnesses.

Surely, these are extreme examples.

Well I think he wanted to show the block when it has zero PE and max KE (the middle picture), and the block when it has max PE and zero KE (the left and right pictures), which is why he chose these three "extreme" examples.

Without extreme examples we wouldn't be able to see the motion. Also, depending on what block we use (a yard stick for example) these examples would be much less extreme.

Agreed- if he showed us the actual shape changes for this example, we wouldn't be able to tell (or if we could, it would be extremely difficult). I like the extreme cases because they are easy to distinguish.

Yeah, although they aren't that realistic, they are useful for getting the point across and I like the figure. I also like how you apply it to the Figure below.

This might just be me nit-picking, but it bothers me that the blocks on the right and the left are different widths than the block in the center, it makes it hard for me to think that they are the same block.

we are making the assumption of only a single mode oscillation in the system here?

The way it is being held wipes out most of the other possibilities (any mode that doesn't have a node where your finger is will get damped out as it tries to move your finger, which is squishy). The surviving modes (in what is called a "pinned-pinned beam") all turn out to have frequencies that are integer multiples of the fundamental. So you hear them as coloration ("timbre" is the official word) on the fundamental. Timbre (the overtone amplitudes) is the difference between a xylophone sound and a piano playing the same fundamental frequency.

I'm surprised that the wood can bend this much.

I don't think it can. It's just to really accentuate what is occurring in the wood.

This are good pictures. Coupled with the diagram of the lattice below of the wood atoms, it is definitely a nice aid to the reading (though the reading is fairly straightforward here).

chemical bonds, which are springs. As the block vibrates, it takes these shapes (shown in a side view):

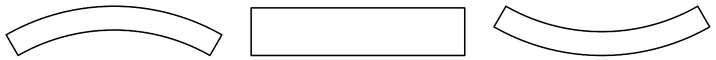

The block is made of springs, and it acts like a big spring. The middle position is the equilibrium position, when the block has zero potential energy and maximum kinetic energy. The potential energy is stored in stretching and compressing the bonds. Imagine deforming the block into a shape like the first shape:

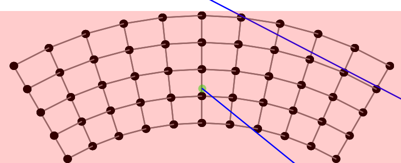

Each dot is a wood 'atom', and each gray line is a spring that models the chemical bond between wood atoms. Deforming the block stretches and compresses these bonds. These numerous individual springs combine to make the block behave like a large spring. Because the block is a big spring, the energy required to to produce a vertical deflection is proportional to the square of the deflection:

$$E \sim ky^2, \tag{9.20}$$

where y is the deflection, and k is the stiffness of the block.

Intuitively, the thicker the block, the stiffer it is (higher k). The spring model will help us find how k depends on the thickness h. To do so, imagine deflecting the thin and thick blocks by the same distance y, then compare their stored energies E_{thin} and E_{thick} by forming their ratio

$$\frac{E_{\text{thick}}}{E_{\text{thin}}}. \tag{9.21}$$

That ratio is

$$\frac{k_{\text{thick}}y^2}{k_{\text{thin}}y^2} = \frac{k_{\text{thick}}}{k_{\text{thin}}} \tag{9.22}$$

because y is the same for the thick and thin blocks. So, the ratio of stored energies is also the ratio of stiffnesses.

The "middle" just refers to where you hit it right?

I think the middle refers literally to the middle, the position that doesn't change during vibration.

I thought he was referring to the middle picture...

Yeah, definitely the picture. Although I guess he could make it less confusing.

Yeah I'm also thought that this was referring to the middle picture in the Figure. I think it has to be since a single point on the block can't be when the block itself has zero potential energy. It has to be a state of the block,which is unbent, to be described like that.

perhaps "neutral" is better terminology here.

with a very very high spring constant

Yes, this seems like a stretch of a comparison. Stiff board = spring..?

doesn't kinetic energy imply motion though?

This section is very well written. I had no problem following the explanation in the paragraphs and the figures add a lot to the reading.

So far this section has been very explicit and easy to follow...due to the pictures and clear emphasis on what is being tested.

Yeah! And it's been especially easy to follow because we saw this example in class.

will you touch on simple beam theory at all in this example and compare it to the use of springs?

This has been the most helpful diagram and explanation for me in a long time.

so then maybe more of the energy is transferred to the thicker block given the sounds from it a higher frequency? If you could hold the blocks more rigidly would the thicker block then have a lower frequency?

chemical bonds, which are springs. As the block vibrates, it takes these shapes (shown in a side view):

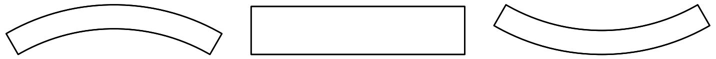

The block is made of springs, and it acts like a big spring. The middle position is the equilibrium position, when the block has zero potential energy and maximum kinetic energy. The potential energy is stored in stretching and compressing the bonds. Imagine deforming the block into a shape like the first shape:

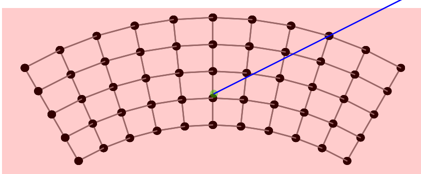

Each dot is a wood 'atom', and each gray line is a spring that models the chemical bond between wood atoms. Deforming the block stretches and compresses these bonds. These numerous individual springs combine to make the block behave like a large spring. Because the block is a big spring, the energy required to to produce a vertical deflection is proportional to the square of the deflection:

$$E \sim ky^2, \tag{9.20}$$

where y is the deflection, and k is the stiffness of the block.

Intuitively, the thicker the block, the stiffer it is (higher k). The spring model will help us find how k depends on the thickness h. To do so, imagine deflecting the thin and thick blocks by the same distance y, then compare their stored energies E_{thin} and E_{thick} by forming their ratio

$$\frac{E_{\text{thick}}}{E_{\text{thin}}}. \tag{9.21}$$

That ratio is

$$\frac{k_{\text{thick}}y^2}{k_{\text{thin}}y^2} = \frac{k_{\text{thick}}}{k_{\text{thin}}} \tag{9.22}$$

because y is the same for the thick and thin blocks. So, the ratio of stored energies is also the ratio of stiffnesses.

this pic is very helpful to me.

I too, like this drawing, but if you just presented us this and said "figure it out now, using springs", i would be very confused, and have even less intuition about what to do next than when i wasn't given this picture. This system looks very complicated and rather intimidating.

I don't really see how this relates to springs, but as a picture of atoms I can see how this framework would behave

The connections between "atoms" have to be springy for this diagram to work. They can't just be rigid rods.

I think the point is that the chemical bond between each 'wood' atom is the same, but the outer ones are stretched which is why they are springy

I really like this diagram

I found this to be nice as well, it was very simple and helped to explain the problem.

Are there really no other complications we would have to deal with ? are we assuming the atoms have no mass?

Can we really assume this? The springs act in two different dimensions, and as the commenter before me mentioned, atoms do have mass (and the springs would have mass too...).

I think it's fine. Springs in every mechanics course are treated to be massless.

I think it goes back to the fundamental part of the class, which is that you have to assume something to get an answer. I think this is definitely within reason.

The atoms definitely have mass: That's the mass that the springs must move. So, more atoms (thicker block) means it is harder to vibrate. On the other hand (as you find later), the thicker block has a lot more spring stretch (in the bonds), and that effect more than makes up for the extra mass to be moved.

Typo.

chemical bonds, which are springs. As the block vibrates, it takes these shapes (shown in a side view):

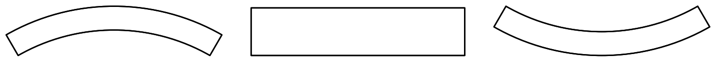

The block is made of springs, and it acts like a big spring. The middle position is the equilibrium position, when the block has zero potential energy and maximum kinetic energy. The potential energy is stored in stretching and compressing the bonds. Imagine deforming the block into a shape like the first shape:

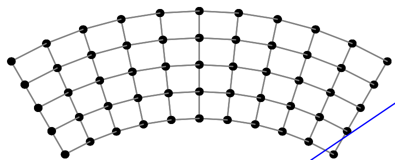

Each dot is a wood 'atom', and each gray line is a spring that models the chemical bond between wood atoms. Deforming the block stretches and compresses these bonds. These numerous individual springs combine to make the block behave like a large spring. Because the block is a big spring, the energy required to to produce a vertical deflection is proportional to the square of the deflection:

$$E \sim ky^2, \tag{9.20}$$

where y is the deflection, and k is the stiffness of the block.

Intuitively, the thicker the block, the stiffer it is (higher k). The spring model will help us find how k depends on the thickness h. To do so, imagine deflecting the thin and thick blocks by the same distance y, then compare their stored energies E_{thin} and E_{thick} by forming their ratio

$$\frac{E_{\text{thick}}}{E_{\text{thin}}}. \tag{9.21}$$

That ratio is

$$\frac{k_{\text{thick}}y^2}{k_{\text{thin}}y^2} = \frac{k_{\text{thick}}}{k_{\text{thin}}} \tag{9.22}$$

because y is the same for the thick and thin blocks. So, the ratio of stored energies is also the ratio of stiffnesses.

Seems like you were describing it as a lot of smaller springs?

Yeah, the chemical bonds between "wood atoms" are springs, as we've established already, so here we're saying that all these individual springs make the combined wood block like one big spring

Is deflection a function of position in this case?

I believe we're looking at deflection of the end of the block, but measuring the deflection anywhere else would only change the constant of proportionality, not the underlying behavior.

Yes, that makes sense. Deflection, or position of a given point in the block relative to equilibrium defines the same relationship with stored energy.

I don't get how we got this

I just compared these variables' dimensions. Since we're talking about springs, I saw the block's stiffness to serve as the spring constant.

[E] = J = M*L^2*T^-2 [k] = M*T^-2 [y] = L

good thinking! This break down may be a good idea to include in the passage for those more unfamiliar with springs or physics in general.

Smart, Springs have stiffness, and so does wood. Therefore, their frequency = sqrt(k/m). Now I understand why the thicker (aka stiffer) wood has higher frequency and thus higher pitch. This is one of thecoolest things I have learned

I wouldn't have thought of this on my own, but after seeing it, it makes a lot of sense.

Same here, it makes so much sense and agreed with intuition but was a little past where I was able to explore the problem.

chemical bonds, which are springs. As the block vibrates, it takes these shapes (shown in a side view):

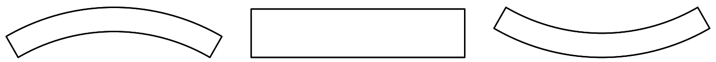

The block is made of springs, and it acts like a big spring. The middle position is the equilibrium position, when the block has zero potential energy and maximum kinetic energy. The potential energy is stored in stretching and compressing the bonds. Imagine deforming the block into a shape like the first shape:

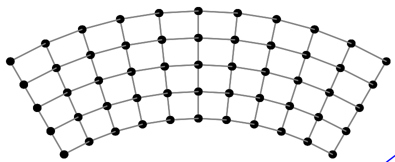

Each dot is a wood 'atom', and each gray line is a spring that models the chemical bond between wood atoms. Deforming the block stretches and compresses these bonds. These numerous individual springs combine to make the block behave like a large spring. Because the block is a big spring, the energy required to to produce a vertical deflection is proportional to the square of the deflection:

$$E \sim ky^2, \tag{9.20}$$

where y is the deflection, and k is the stiffness of the block.

Intuitively, the thicker the block, the stiffer it is (higher k). The spring model will help us find how k depends on the thickness h. To do so, imagine deflecting the thin and thick blocks by the same distance y, then compare their stored energies E_{thin} and E_{thick} by forming their ratio

$$\frac{E_{\text{thick}}}{E_{\text{thin}}}. \tag{9.21}$$

That ratio is

$$\frac{k_{\text{thick}}y^2}{k_{\text{thin}}y^2} = \frac{k_{\text{thick}}}{k_{\text{thin}}} \tag{9.22}$$

because y is the same for the thick and thin blocks. So, the ratio of stored energies is also the ratio of stiffnesses.

Is this like U = .5kx^2 for normal springs?

I am confused here; why not having it relate to a deflection angle? Depending on where you are, the deflection is different, no?

I'm thinking of the block as a black box (an abstraction!). You push on it, and it pushes back (like a spring would). The question is how hard it pushes back for a given deflection, however that deflection is measured. The "how hard per deflection" is the spring constant.

(To answer the original post: Right, the ky^2 is just the U=kx^2/2 for a spring, but dropping the factor of 2 because 1=2 in this course.)

maybe you should explain how to test for stiffness. if you think about testing for stiffness, its related to how much deflection you see compared to the force you are applying.....and you can measure force applied on a spring scale. so yes, it is like hookes law, right?

Could you maybe explain the derivation of this equation? or origin?

I think of k as a youngs modulus of sorts- and that does not change with thickness it independent of size and shape

isn't EACH of the k's for each of the springs the same? b/c it is the same material. there are just more rows of k's, so the effective k is higher.

That sound right, but he's just talking about the "combined" spring, or the 'effective' spring. If we think of the whole block as one spring, then it has only 1 k, which just happens to be a combination of the individual ks.

I agree, think about what happens on the macro scale when you add springs in parallel, the effective stiffness increases.

Also it makes sense that the thickness of the wood is the dimension we care about since the "springs" stretched during perturbations are the horizontal ones. So thicker wood means more horizontal springs.

to the original commenter, you're assuming k is an intensive property, which it isn't. (It's not an extensive property either – in fact, it works fairly similarly to conductance, the inverse of resistance.)

chemical bonds, which are springs. As the block vibrates, it takes these shapes (shown in a side view):

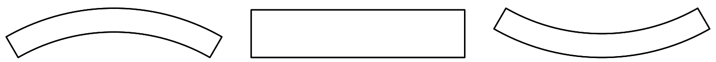

The block is made of springs, and it acts like a big spring. The middle position is the equilibrium position, when the block has zero potential energy and maximum kinetic energy. The potential energy is stored in stretching and compressing the bonds. Imagine deforming the block into a shape like the first shape:

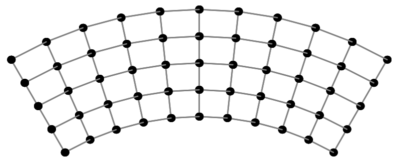

Each dot is a wood 'atom', and each gray line is a spring that models the chemical bond between wood atoms. Deforming the block stretches and compresses these bonds. These numerous individual springs combine to make the block behave like a large spring. Because the block is a big spring, the energy required to to produce a vertical deflection is proportional to the square of the deflection:

$$E \sim ky^2, \tag{9.20}$$

where y is the deflection, and k is the stiffness of the block.

Intuitively, the thicker the block, the stiffer it is (higher k). The spring model will help us find how k depends on the thickness h. To do so, imagine deflecting the thin and thick blocks by the same distance y, then compare their stored energies E_{thin} and E_{thick} by forming their ratio

$$\frac{E_{thick}}{E_{thin}}. \tag{9.21}$$

That ratio is

$$\frac{k_{thick}y^2}{k_{thin}y^2} = \frac{k_{thick}}{k_{thin}} \tag{9.22}$$

because y is the same for the thick and thin blocks. So, the ratio of stored energies is also the ratio of stiffnesses.

How do we know that tapping the blocks leads to the same deflection?

> It doesn't matter! For two reasons. First, one can tap them in such a way that the deflections are the same. Second, even if that isn't true, frequency is independent of amplitude (the key feature of simple harmonic motion). So, we can assume any amplitude we want and get a frequency, and it'll be the same for all amplitudes. Every time I run across this SHM property, I am surprised.

This is such a better example then the planets in the last reading

This section is very clear but I don't know if it follows with the rest of the text. There are much more intense assumptions that are made in the rest of the readings while this is a very simple thing to understand and is very explicitly stated.

I think that's good though. This is explaining and illustrating how spring models work. It wouldn't help much if we were confused while reading it.

why is the deflection the same for the thin and thick blocks? Is it that the amplitude is determined by the input force?

He said earlier that we're going to deflect the blocks the same amount.

> It doesn't matter! For two reasons. First, one can tap them in such a way that the deflections are the same. Second, even if that isn't true, frequency is independent of amplitude (the key feature of simple harmonic motion). So, we can assume any amplitude we want and get a frequency, and it'll be the same for all amplitudes. Every time I run across this SHM property, I am surprised.

So energy is proportional to pitch? Higher energy -> higher pitch?

Yes, high energy <-> high frequency/pitch

and we can just ignore the added mass and other things addressed in the beginning? actuallly i guess he didnt ask us to predict how much the frequency would CHANGE, just which would be higher

This is a really good example - very clear and interesting. When you mentioned the ratio of energies, a lightbulb went off in my head, and it all made perfect sense.

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

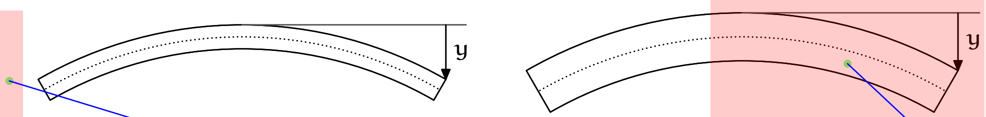

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \qquad (9.23)$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \qquad (9.24)$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \qquad (9.25)$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \qquad (9.26)$$

It's so interesting that this is your favorite demonstration.

then why is the dotted line bent??

it's that the material below the line is in compression, and the material above it is in tension. the neutral line is in neither compression or tension.

This is off topic- but these diagrams remind me of the pictures of illusions usually shown in psych classes...the one that comes to mind is the picture of 2 lines that are the same length but look different because of the direction of the arrows on each end (<——> and >——<).. for this diagram (at least to me), the y on the right looks bigger than the y on the left because of the different in the blocks, but they are actually the same!

this is a vry good picture to help understand what is happening

It might be nice to put the same diagram with the bond lines and dots here to go with this text. These two diagrams don't illustrate the concept as well.

i'm having trouble picturing this

Is he referring to the "neutral line"?

I think it would have been better to see the entire motion of each spring. It would make the explanation a little clearer.

i think it refers to the bonds across the thickness of the plank rather than the length (aka the 'radial' direction.)

I feel like just tapping the block will not lead to a measurable deflection in the beam

I think you could explain a little more here... I am starting to get lost.

I don't really understand this scenario. Which springs are the radial direction springs? maybe detailing this in one of the figures would be helpful

The bent wood appears to be a piece of a circle (an arc), and the radially directed springs would be the ones along the radius of this circle. These springs would only come into play if the thickness of the board changed along its length.

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

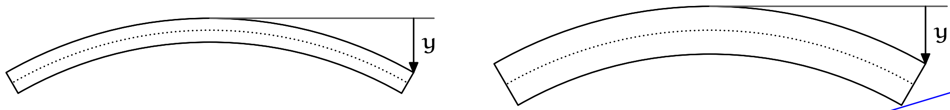

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

I'm kind of confused as to what we are comparing.

We're looking at how far our "springs" are stretched in each block for the same deflection

I feel like this section only makes sense to me because it reminds me of beam bending in 2.001, but otherwise I'd be confused since there is no diagram...

yeah I haven't taken 2.001 and I'm confused about this section..

I haven't taken 2.001, but it seemed pretty clear to me what this paragraph is talking about based on the previous diagram.

I'm a bit lost here...

I feel like this whole section could be a lot more easily understood with a few simple diagrams of the differences between the two blocks at the molecular level, instead of the existing images, which show only the macroscopic structural differences.

I agree, it would be nice to see the springs shows in the pictures as opposed to these images.

Definitely. And perhaps bullet points to show where each factor of 2 comes from, instead of hiding them in the paragraph.

Yeah I agree with the factors of two comment. I honestly came out of the paragraph not understanding where the 8 was coming from. I had to re-read the paragraph.

I agree with all the comments in this thread and will add many diagrams as a result. (My only defense is that you should have seen the reading in its state a few days ago, before I added many of the diagrams and improved other diagrams – because I knew it would be helpful based on earlier memos).

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

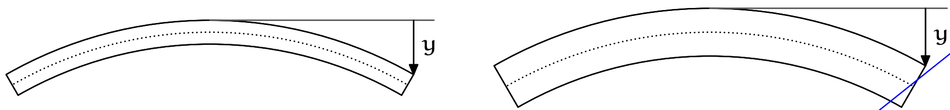

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

Why would the spring in the thick block have twice the extension of a spring in the thin block? I would have guessed it's the other way around

I am also very confused by this.

This is confusing, could you include a diagram and point this stuff out? I think that could help a lot.

Take the case of a spring bond at the top edge of a block: This spring in the thick block is twice the distance from the neutral axis as that in the thin block. As he notes, the amount of extension is proportional to the distance from the neutral axis. This fact comes from looking at the radius of curvature and strain in the block. A decent derivation: <http://www.ecourses.ou.edu/cgi-bin/ebook.cgi?doc=&topic=me&chap_sec=04.1&am

Ah, that makes sense! It would help to have that explained more explicitly. I feel like this section so far was good about going slowly and then this paragraph just kind of whirled by...

The thick block has twice as many "wood" atoms, so the springs should extend the same since there twice as many to span the doubled extension, no?

True there are more wood atoms, but in term of matching up atoms from the thin block with the thick block, atoms would be like "every other" one, hence twice as far.

ok, but how does that relate to 4 times the energy?

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

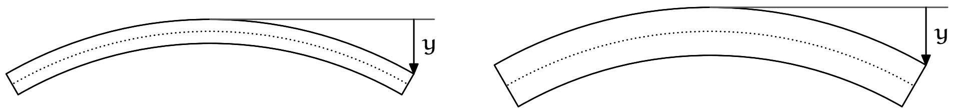

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

This point definitely needs a few diagrams to make it more intuitive (as suggested in a few other threads). Imagine pairing each spring in the thin block with a spring in the thick block that is twice as far from the neutral line. So, the neutral line maps to the neutral line; the outer edge maps to the outer edge (and same for the inner edge).

The next idea is that the spring's extension is proportional to distance from the neutral line. Think of each row as an arc of a circle. As you go outward from the neutral line, the radius of the arc increases linearly, as does the arc's length. But all the arcs have the same number of springs (or atoms) – that is determined by the unstretched length. So, the extra stretch, which is proportional to distance from the neutral line, is distributed over a fixed number of springs. Therefore, each spring's extension is proportional to the distance from the neutral line. (Which means that the energy stored in each spring is proportional to the square of the distance from the neutral line.)

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

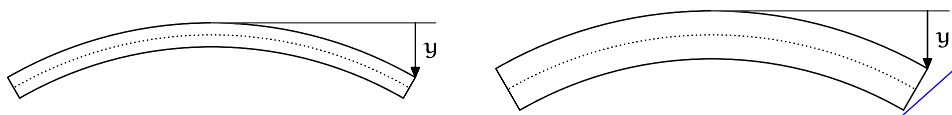

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

i don't have this entirely clear in my head yet, but i think there's a second moment of inertia that has something to do with this.

Yes, this paragraph goes by much too fast. I think you could add some information to the picture above to make it more clear, perhaps. I'm not entirely sure, for example, why twice the extension leads to 4 times the energy.

The energy stored in a spring is 1/2*K*xˆ2. If you displace it by twice some reference displacement, you get (2x)ˆ2 –> 4xˆ2, which when plugged in to the energy formula gets four times the reference energy.

I think it would be helpful to have that formula for the energy stored in a spring in the actual text for those students who have forgotten some 8.01.

I agree, this paragraph needs some diagrams or more step by step equations showing relationships between quantities/proportionality

I agree about the spring energy storage, for non course 2 people this might not be information readily stored in one's head

Yeah, this paragraph was pretty confusing for a non-course 2 person. Specifically I'm confused on the jump from 4 to 8 times more energy... could you explain the bit about each spring has 2 partners?

If both springs are displaced by the same amount as shown in the picture above, why would there be a factor of 8?

I disagree. This was very clear. It was definitely a 'wow' moment. Very good explanation.

This strikes me as a little hand-wavey. I get it, I think, but it wouldn't necessarily make sense if I hadn't taken 2.001

it feels like the distance and the double layers is counting the same thing twice. can we go over how they are different?

should be "block"

all the 'thin' and 'thick' words are confusing me

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

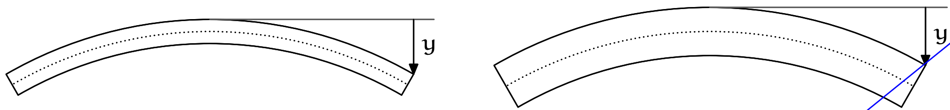

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

very well explained

It took me a while, but I think I understand it. It would be nice to somehow demonstrate this using the above images.

perhaps you can use this space to also specify what 2s were multiplied since they're stated in the above and below paragraph but not summarized.

This sentence is misleading. I understood the paragraph above it, but this sentence threw me off for a second. The ratio of the thicknesses is 2, and you say to multiply by 2, which would lead me to a conclusion of 4, which is, of course, not 8.

i think you mean for a piece twice as thick, rather than "multiplying"

In what units is 'stiffness' measured?

Force per displacement (so, mass*length/timeˆ2 divided by length => mass/timeˆ2).

stiffness in this case is k, the "spring constant"

Would it be better to mention this fact before jumping into the ratio equating to 8 and then explaining later?

This answers where the 8 came from. The above paragraphs made the reader want to guess at how it came from before you arrive at the explanation more concisely. Maybe a reordering of the timing of when you present this would be good?

I think this is just a good thing to memorize. Does it only hold for wood?

it sounds like it can be more general ("in general"), though I have never heard this before

As long as you're going to memorize, just remember k b*hˆ3 (i.e. first power in beam width, third in beam height). This comes from the bending moment of inertia of a rectangular beam, which is I=b*hˆ3/12, and k YI (Y is young's modulus)

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

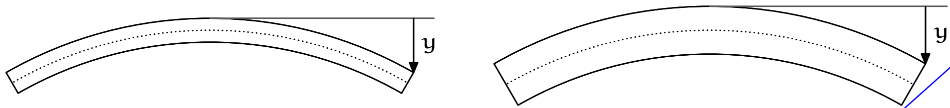

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

And all this time, I'd always assumed k stiffness to be an intensive property. Goes to show you even the basest assumptions should be double-checked.

One example that I often think of to remind myself that the stiffness isn't intensive is that thin sheets of metal (or any other substance, for that matter) are much easier to transform than thick blocks. Therefore, the stiffness must increase with thickness.

Now I remember, the *elastic modulus* is the intrinsic property! THAT'S how you can derive the stiffness with other knowns. I'm really glad I remembered that, I think I would've gone crazy otherwise.

I can follow how the bonds between atoms act like springs, but how is the entire wood block like a spring now?

If you get two springs, and connect them, you still have a spring! Now, if you have this elaborate network of springs, at the end, you still get a spring out.

You know how complicated circuits with resistors have a resistance, it's kind of like that. All these elaborate springs in the end have some effective "k".

I can follow how the bonds between atoms act like springs, but how is the entire wood block like a spring now?

The entire wood block is made up of these springs therefore the entire thing acts like a spring. Imagine a set of springs all hooked up to one another.

Yeah, if you remember in mechanics, they dealt with problems where two or more springs are connected together... then we solved to find that it acts like a "new" spring with a "new" spring constant that is expressed as a function of the individual spring constants. I guess springs preserve their spring-like nature from a microscopic level to a macroscopic level, sort of like fractals.

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

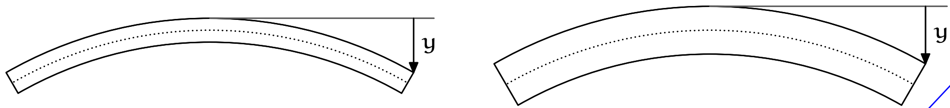

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

This is a good formula to remember

I agree, it pops up all over the place in mechanical physics because it works for all different "effective stiffnesses" and "effective masses."

Are we expected to know this off the top of our heads? I'm course 6, and I certainly don't know this formula.

Actually, you can derive it pretty quickly by dimensional analysis. Omega has units of 1/time, and the only other given quantities are k (units of force/length = mass/timeˆ2) and m (units of mass), so to get units of 1/time, we must take the square root of k/m. It also makes intuitive sense because increasing k (the stiffness) will make the block return to its original shape faster (higher frequency), while increasing m will make the block harder to accelerate (lower frequency).

Yeah I didn't know this one either, it seemed like a very convenient equation without a reason. Very good explanation of unit analysis for this though.

yeah it's very handy to know the natural frequency of materials, especially if you're designing something that will move. you want to make sure that the part's natural frequency is a lot higher than the vibrations in the piece so it doesn't cause resonance.

You also get exactly the same issues designing certain kinds of circuits – you want your switching frequency to be far above or below the resonant frequency of the circuit (unless resonance is your goal, of course). It's pretty cool to see the ways that we're all using the same basic concepts...

I think its incredible that we can get to this point just by reasoning, you just have to think about things the right way

Definitely...this example is awesome simply because the conclusion is so concrete and correct, despite the beginning of the section clearly outlining the difficulty in making the guess whether or no the ratio will be greater or less than 1.

I think I'll use this example as the first physical example of springs. Maybe I'll use the cosine integral as a short introduction to the parabolic approximation in general, then explain how the interatomic bond potentials have the same property – which means they are springs. Then ask the wood-blocks question and analyze it. (The planets should probably come later – or never.)

To find the stored energies, look at this picture of the blocks, with the dotted line showing the neutral line (the line without compression or extension):

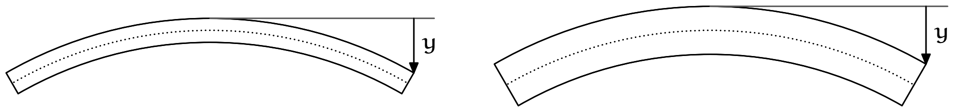

The deflection hardly changes the lengths of the radial-direction bond springs. However, the tangential springs (along the long length of the block) get extended or compressed. Above the neutral line the springs are extended. Below the neutral line, the springs are compressed. The amount of extension is proportional to the distance from the neutral line.

Now study comparable bond springs in the thin and thick blocks. Each spring in the thin block corresponds to a spring in the thick block that is twice as far away from the neutral line. The spring in the thick block has twice the extension (or compression) of its partner in the thin block. So the spring in the thick block stores four times the energy of its partner spring in the thin block. Furthermore, the thick block has twice as many layers as does the thin black, so each spring in the thin block has two partners, with identical extension, in the thick block. So the thick block stores eight times the energy of the thin block (for the same deflection y).

Thus

$$\frac{k_{\text{thick}}}{k_{\text{thin}}} = 8. \tag{9.23}$$

This factor of 8 results from multiplying the thickness by 2. In general, stiffness is proportional to the cube of the thickness:

$$k \propto h^3. \tag{9.24}$$

Because the entire wood block acts like a spring, its oscillation frequency is $\omega = \sqrt{k/m}$. The mass ratio is caused by the thickness ratio:

$$\frac{m_{\text{thick}}}{m_{\text{thin}}} = 2. \tag{9.25}$$

Because the stiffness ratio is 8, the frequency ratio is

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{8}{2}} = 2. \tag{9.26}$$

This reasoning makes a lot of sense to me.

This actually took me a couple minutes to figure out. I somehow overlooked the w=(k/m)^-.5 in the previous line. Perhaps that would be better placed right before this ratio of w's.

Yeah this little math reasoning right here is obvious to me, but I never would have come up with the rest..and its pretty crazy that frequency turns out to be directly proportional to thickness

so if it were 8 times as thick then the frequency ratio would be 8? pretty neat. Does this scale for any thickness?

There must also be some point where the object simply doesn't produce any audible sound. The volume of the sound must decrease with thickness, or the object simply becomes so stiff that hitting it only locally deforms the object instead of making the whole thing oscillate, it would be interesting to see where this point lies, and see how the frequency really does/doesn't scale with frequency.

That makes sense since volume is proportional to how much the object perturbs. So a less stiff plank would move more air around making it louder.

In general, $m \propto h$ so

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{h^3}{h}} = h. \qquad (9.27)$$

Frequency is proportional to thickness!

Comments on page 4

I think you should give a overall summary of this relevance. Connect it back to the beginning of the chapter, and site why this is important.

I really liked this analysis- Thanks!

so k is also prop. to h?

k is proportional to hˆ3

This was a cool analysis, but I'm wondering if we didn't use the fact that we knew what happened to explain why it happened. To me, it's cooler to go from basic principles to having a prediction. I think we did this in this section–but it might be more impactful to have the analysis first, then the experiment to confirm our analysis.

although, it's useful to note that science tends to work in the opposite way: someone notices something and tries to explain the phenomenon with the analysis.

Cool!

This lecture seems relatively simple, but I still am not sure if I completely understand it.

is this assuming the other dimensions are effectively infinitely larger?

or that the thickness is very thin.

Yes, it does, or at least within a large enough range.

Okay. Thanks for the clarification. I originally missed this distinction.

Are there any types of materials for which this doesn't apply?

Wouldn't this still apply, regardless of the thickness? You just might need a larger force to excite the block to an audible frequency of oscillation...

I would never have thought that. I was sure it would have been the opposite.

Yea, I assumed it would be proportional to thickness but in the sense that the thicker wood would change the wavelength, not in the way of affecting the vibration of the wood

yeah this is a really cool result, that I would not have expected!!

In general, $m \propto h$ so

$$\frac{\omega_{\text{thick}}}{\omega_{\text{thin}}} = \sqrt{\frac{h^3}{h}} = h. \tag{9.27}$$

Frequency is proportional to thickness!

How did this compare to the results from your experiment?

It would be nice to mention this (some numbers/data would do). I remember in class you said that the frequency of the thick block was nearly an octave higher than the frequency of the thin block.

I agree. In lecture you mentioned it was actually a Major 6th interval. Is the discrepancy from geometry or material error, edge effects, or something else (like something more subtle that's affecting the frequency)?

So how does this result relate to the xylophone problem from the pretest (where we are adjusting length not thickness)?

This finding is a different relationship. From what I remember, in the xylophone example, shorter bars corresponded to lower pitches and longer bars to higher pitches.

I thought is was the same relationship. Both examples deal with 2 constant dimensions and 1 dimension that varies. The only difference between the xylophone and the block-example is which dimension varies between multiple planks.

Ah, that's the subject of lecture today (I have a xylophone)!

I'm sad I missed that! And I'm pretty sure that the longer bars are lower...

Might be nice to tie this conclusion into the first paragraph or two...

I agree, the most surprising part to me is that it is linearly proportional to the thickness. A little recap would have brought things together well.

Agreed, just a short recap would be great. It's not absolutely necessary, but it's always nice just to get a few sentences tying up the issue or problem that was presented a few pages ago. Even just "the thicker block has a higher frequency" would be great.

Wow, this was pretty interesting that this could all be explained using 8.01 concepts.

Interesting example! I liked that it was less mathy in the reasoning and relied a lot on just thinking about the picture also.

Is this the end of the textbook?