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 5ˆ2 * 3 (for pi), you get 75. 10ˆ2 is 100. The actual is 6ˆ2 *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*rˆ2. 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 $(10\,\mathrm{cm})^2$ 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 $(10\,\mathrm{cm})^2$ 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 $(10\,\mathrm{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

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

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 nerddom 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 $(10\,\mathrm{cm})^2$ 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 $(10\,\mathrm{cm})^2$ 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 $(10\,\mathrm{cm})^2$ area. 10^{10} pits would need 10^5 rows and 10^5 columns, so the spacing between pits is roughly

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

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 10ˆ10 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(10ˆ10/2)=10ˆ5/sqrt(2) rows, so d 10 cm/(10ˆ5/sqrt(2)) = 10 cm / 10ˆ5 *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 10ˆ10 pits would require 10ˆ5 rows and 10ˆ5 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*10ˆ9 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 $3600\,\text{s}/1\,\text{hr}$; in other words, $3600\,\text{s} = 1\,\text{hr}$. 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, $3600\,\text{s} = 1\,\text{hr}$. 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:

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

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.

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 $3.6 \times 4.4 \times 3.2$. 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

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 $3.6 \times 4.4 \times 3.2$. 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

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.