# 1 Divide and conquer

Faced with the difficult problem of ruling vast territory filled with hostile groups, empires found the same solution: Set the groups fighting one another rather than the empire. Using the language of one of the technique's masters: *divide et imperia* or, in modern language, divide and rule. We illustrate the value of this imperial advice by applying it to a few problems.

We use everyday estimations. They warm you up for more complicated problems and illustrate how to sanity check information that you see. Information abounds; newspapers put out billions of words a year (or is it trillions?). So the problem of our time is not in finding information but in evaluating it.

---

*Health care*

Suppose a newspaper article asserts that the annual cost of health care in the United States will soon surpass $1 billion. Does this number seem reasonable, is it far too small, or is it far too large?

---

## 1.1 Compact discs

*How far are the pits spaced on a compact disc (CD)?*

Each pit stores one binary digit (bit) of information, so a CD (not drawn to scale) looks like

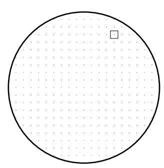

The small square indicates the zone governed by one pit, and the length of the square's side is an estimate of the spacing between pits. With a simpler shape the spacing is easier to estimate:

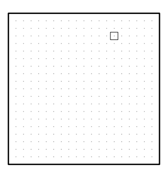

From the area of the CD and from how many bits it stores (its capacity), we can find the spacing. This **tree diagram** represents the information that the spacing depends only on the area and the capacity:

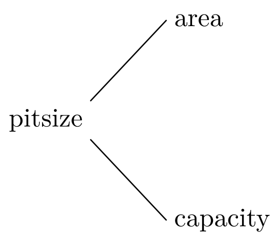


However, our work is not done. The endpoints of the tree, the leaf nodes, are not simple enough. First, areas are hard to estimate directly. Most of us do not have a ready mental library of areas to compare against. For example, how large is $100\,\mathrm{m}^2$? Oh, you say, that's no problem: It's about 10 m tall and 10 m wide. This answer reiterates the point: An area is hard to perceive directly, so we divide it into two lengths because we have more experience judging lengths than judging area. The tree springs two branches:

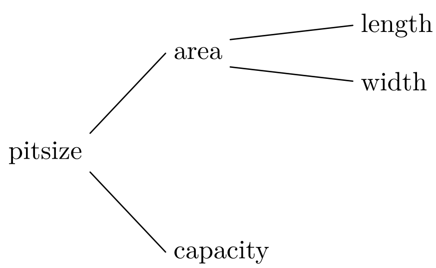


Following that procedure for the CD, the area is

$$A \sim 10\,\mathrm{cm} \times 10\,\mathrm{cm}. \tag{1.1}$$

The equation introduces the ubiquitous $\sim$ symbol known to its friends as **twiddle**. It indicates that, unlike with the $\propto$ symbol ('proportional to'), the two sides of the equation have the same dimensions – in this case, area – but that the equation may be missing a numerical factor like 5 or 10. Another example of twiddle is

$$\text{kinetic energy} \sim Mv^2, \tag{1.2}$$

which is missing a factor of one-half on the right side. Using twiddle is healthy. One reason to use divide-and-conquer reasoning is that it helps you start solving a problem. But what use is starting if you soon bog down deciding the exact dimensions and shape of the CD? The twiddle, an indulgence to commit numerical sins, frees you from that trap. So

$$A \sim 10\,\text{cm} \times 10\,\text{cm} = 100\,\text{cm}^2. \tag{1.3}$$

The capacity is also not obvious. One tactic is to look at a data CD. Those claim a capacity of roughly 700 MB. But let's not accept that value on faith. Instead, break it into pieces. The CD is designed to store music represented by frequent digital samples of the loudness. That history suggests breaking the capacity into the duration, the sample size (bits per sample), and the sampling rate:

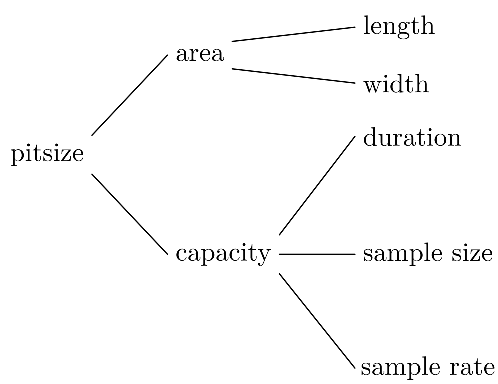


For the duration, a CD holds about an hour of music: one long classical symphony or 20 popular songs that are about 3 minutes each. The pop-song method turns the duration into two leaves:

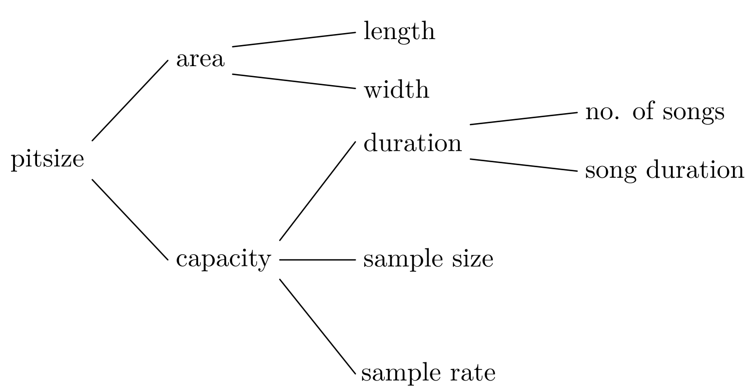

The unevaluated leaves are the sampling rate and the sample size. For the sample size, remember that the sound reproduction is designed to be high fidelity, so each sample should have a decent number of bits, perhaps 16. And the CD reproduces sound in stereo, which requires two channels, making 32 bits per sample. This tree reflects breaking up the sample size:

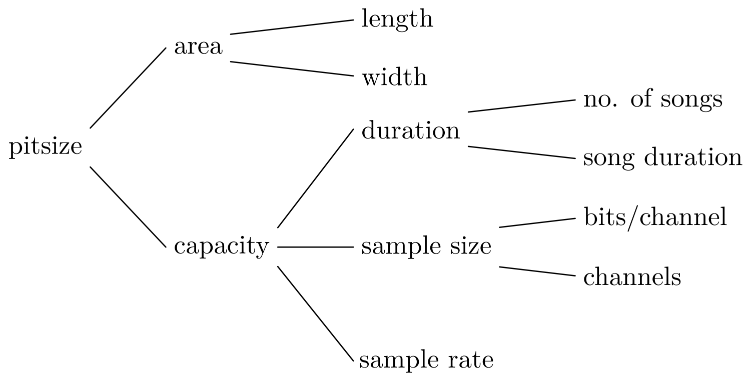


Only the sampling rate remains to evaluate or break apart. The Nyquist sampling theorem says that the sampling rate must be at least twice the highest frequency in the signal, which breaks up the sample rate:

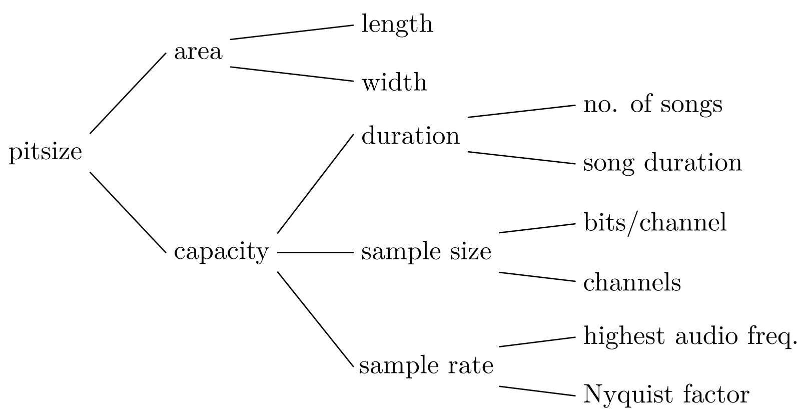


Since the highest frequency that humans can hear is around 20 kHz, the required sampling rate is $\sim 40\,\text{kHz}$.

Now gather the numbers that make the result. The capacity is

$$N \sim 3600\,\text{s} \times \frac{4 \cdot 10^4\,\text{samples}}{1\,\text{s}} \times \frac{32\,\text{bits}}{1\,\text{sample}} \sim 5 \cdot 10^9\,\text{bits}. \qquad (1.4)$$

These bits are scattered over an area of $100\,\text{cm}^2$, so each bit occupies an area

$$a \sim \frac{100\,\mathrm{cm}^2}{5 \cdot 10^9}. \qquad (1.5)$$

The division would be easy if the denominator were a pure power of 10. If you multiply it by 2 then it would become $10^{10}$. To undo the damage from that change, multiply the numerator by 2 as well. So think of the calculation like

$$a \sim \frac{100\,\mathrm{cm}^2}{5 \cdot 10^9} \times \frac{2}{2} = \frac{2 \times 100\,\mathrm{cm}^2}{10^{10}} =\sim 2 \cdot 10^{-8}\,\mathrm{cm}^2. \qquad (1.6)$$

A square of this tiny area has a side length

$$x = \sqrt{a} \sim \sqrt{2 \cdot 10^{-8}\,\mathrm{cm}^2}, \qquad (1.7)$$

which is roughly $10^{-4}$ cm or 1 micron. This length is comparable to the wavelength of light, so the CD will act like a diffraction grating. Ah, that's why the CD shimmers with beautiful reds, blues, greens, and colors in between!

The point of this first example is to teach you the habit of making trees as you break a problem into ever smaller parts. The next example, which provides more practice with trees, we dedicate to physicists who need employment outside of physics.

## 1.2 Bills or gold?

*Having broken into a bank vault, should you take the $100 bills or the gold?*

The choice depends on how easily and losslessly you can fence the loot and on other issues outside the scope of this book. But we can study one question: Which choice lets you carry away the most money? The weight or the volume may limit how much you can carry and, more importantly for this problem, affect your choice. To make a start, let's assume that you are limited by the weight (actually, the mass) that you can carry, The problem then depends on two subproblems: the value per mass for $100 bills and for gold. In tree form:

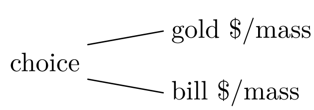


The value per mass of gold might be a familiar figure from the newspaper or from the financial section of the evening news. It's now (2006) about $600/oz (oz being the abbreviation for an ounce). As a rough check on the memory –

e.g. should it be \$60/oz or \$6000/oz? – here is another method. When the gold standard was reintroduced as the dollar standard in 1945, gold was set at \$35/oz. Inflation has probably devalued the dollar by at least a factor of 10 since then, so gold should be at least \$350/oz now. The half-remembered figure of \$600/oz seems reasonable.

Finding the value per mass of a dollar bill requires subdivision:

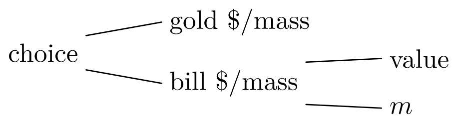


The value is given (\$100) but the mass needs work. It breaks into the volume times the density:

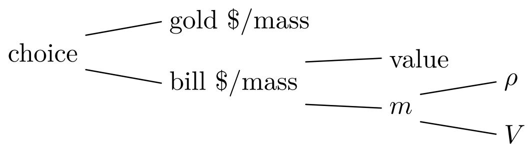


The volume itself breaks into length times width times thickness:

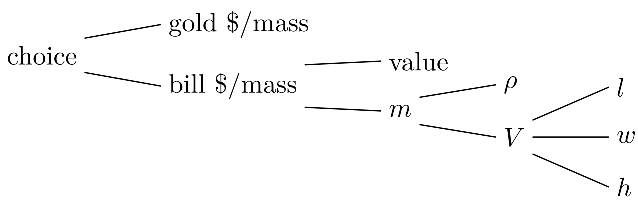


For the length and width, lay a ruler next to a dollar bill or guess that a bill measures 2 or 3 inches by 6 inches or $6\,\mathrm{cm} \times 15\,\mathrm{cm}$. To develop your feel for sizes, guess first; if you feel uneasy, check your answer with a ruler. As your feel for sizes develops, you will need the ruler less frequently.

The thickness is harder. We take any piece of information that we can get. As George Washington Plunkitt, onetime boss of Tammany Hall, said (Riordan, 1963, p. 3): ‘I seen my opportunities and I took ’em.’ We pretend that a dollar bill is made from ordinary paper. To find its thickness, look around. Next to the computer used to compose this textbook sits an inkjet printer; next to the printer is a ream of printer paper. If we know how thick the ream is and how many sheets it has, then we know the thickness of one sheet. You might call this technique multiply and conquer. Jesting aside, the general lesson is that tiny values, those much below typical human experience, need to be magnified to make them easy to estimate. Large values, those much above typical human experience, need to be broken into smaller parts to make them easy to estimate. With this last step of magnifying the sheet's thickness, the full tree becomes

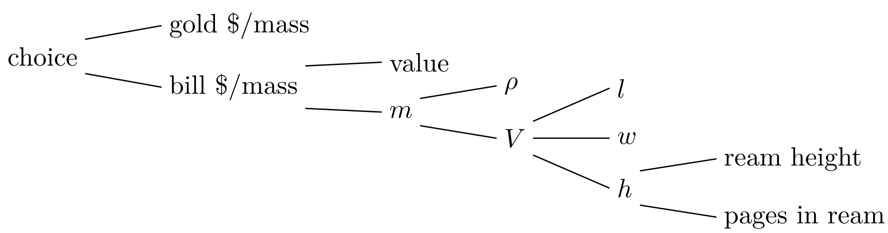


The ream (500 sheets) is roughly 5 cm thick, so a sheet of quality paper has thickness $10^{-2}$ cm. The only missing value is the density of the bill. To find it, use what you know. Money is paper. Paper is wood or fabric, except for many complex processing stages whose analysis is beyond the scope of this book. When a process, here papermaking, looks formidable, forget about it and hope that you'll be okay anyway. More important is to get an estimate; you can correct the egregiously inaccurate assumptions later. How dense is wood? Wood barely floats, so its density is roughly that of water: $\rho \sim 1\,\mathrm{g\,cm^{-3}}$.

Now we have assembled the pieces to compute the mass of the bill

$$m \sim 6\,\mathrm{cm} \times 15\,\mathrm{cm} \times 10^{-2}\,\mathrm{cm} \times 1\,\mathrm{g\,cm^{-3}} \sim 1\,\mathrm{g}. \qquad (1.8)$$

So the value per mass of a $100 bill is $100/g. To choose between the bills and gold, compare that value to the value per mass of gold. Unfortunately our figure for gold is in dollars per ounce rather than per gram. Fortunately one ounce is roughly 30 g so $600/oz is roughly $20/g. Moral: Take the bills.

---

*Volume*
How do volume limits affect the decision between gold and bills?

*Armored car*
How much money would you expect to find in an armored car?

---

## 1.3 Why it works

Breaking problems into factors, besides making the estimation possible, often reduces the error in the estimate. If you guess a number of the order of $10^{10}$ in one step, you might be in error by a factor of 10. For example, you might estimate the number of stars in our galaxy as $N \sim 10^{10}$, but $10^9$ or $10^{11}$ might feel equally plausible. Now break the estimate into two pieces:

$$N \sim A \times B, \tag{1.9}$$

where $A$ and $B$ are each of order $10^5$. Now you estimate $A$ and $B$. What is the typical error in your estimate of $N$? There probably is a general rule about guessing, that the *logarithm* of a number is in error by a fixed fraction. If estimates of order $10^{10}$ are often in error by a factor of 10, which is one unit on a log-base-10 scale, then estimates of order $10^5$ would be in error by one-half of a unit or by a factor of 3. In the one-shot estimate for $N$, its logarithm $\log_{10} N$ could be 9, 10, or 11 – the three values feeling equally plausible:

$$\begin{array}{cccc} \log_{10} N & 9 & 10 & 11 \\ p & \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{array} \tag{1.10}$$

For $A$ or $B$, each of order $10^5$, their logarithms could be 4.5, 5, or 5.5, each value feeling equally plausible. Now see what happens when you multiply $A$ and $B$ or, equivalently, when you add their logarithms:

$$\begin{aligned} \log_{10} N &= \log_{10} A + \log_{10} B \\ &= \left\{ \begin{array}{c} 4.5 \\ 5 \\ 5.5 \end{array} \right\} + \left\{ \begin{array}{c} 4.5 \\ 5 \\ 5.5 \end{array} \right\}, \end{aligned} \tag{1.11}$$

where the curly braces list equally plausible values. Since each of $A$ and $B$ have three possibilities, their sum has nine possibilities including duplicates. Here are the possibles sums and their probabilities, assuming that the errors in $A$ and $B$ are uncorrelated:

$$\begin{array}{cccccc} \log_{10} N & 9 & 9.5 & 10 & 10.5 & 11 \\ p & \frac{1}{9} & \frac{2}{9} & \frac{3}{9} & \frac{2}{9} & \frac{1}{9} \end{array} \tag{1.12}$$

Compared to the one-shot distribution, this distribution is more likely to produce a $\log_{10} N$ close to 10. In other words, breaking the estimate into two parts reduces the expected error, because the errors in the parts have a chance to cancel.

Generalizing this example, we break $N$ into $k$ roughly equal parts; the example used $k = 2$. The estimate of each part is then in error by a factor of $\gamma = 10^{1/k}$. If these errors are uncorrelated, their logarithms combine as steps in a random walk. So the error in $\log_{10} N$ arises from a random walk of $k$ steps, each with step size $\log_{10} \gamma = 1/k$. In such a random walk, the expected root-mean-squared distance from the origin is

$$x_{\text{rms}} = (\text{number of steps})^{1/2} \times \text{step size}, \tag{1.13}$$

which is

$$x_{\text{rms}} = k^{1/2} \times \frac{1}{k} = k^{-1/2}. \tag{1.14}$$

This distance from the origin is the typical error in the logarithm of $N$. By increasing $k$, you decrease this distance and therefore decrease the error in $N$. The moral: **Divide and conquer!**

---

*Batteries*
What is the cost of energy from a 9V battery? From a wall socket (the mains)? Draw a tree to illustrate your solution.

*Human warmth*
How much heat do you generate just sitting around?

*Fuel economy*
What is the fuel consumption, in passenger–miles per gallon, of a 747 jumbo jet?

*Bandwidth*
What is the data rate (bits/s) of a 747 filled with DVD's crossing the Atlantic? Draw a tree to illustrate your solution.

---