

Learning Objectives:

- Introduce and motivate the study of computation for its own sake, irrespective of particular implementations.
- The notion of an *algorithm* and some of its history.
- Algorithms as not just *tools*, but also *ways of thinking and understanding*.
- Taste of Big-*O* analysis and surprising creativity in efficient algorithms.

0

Introduction

*“Computer Science is no more about computers than astronomy is about telescopes”, attributed to Edsger Dijkstra.*¹

“Hackers need to understand the theory of computation about as much as painters need to understand paint chemistry.” , Paul Graham 2003.²

“The subject of my talk is perhaps most directly indicated by simply asking two questions: first, is it harder to multiply than to add? and second, why? ...I (would like to) show that there is no algorithm for multiplication computationally as simple as that for addition, and this proves something of a stumbling block.”, Alan Cobham, 1964

¹ This quote is typically read as disparaging the importance of actual physical computers in Computer Science, but note that telescopes are absolutely essential to astronomy as they provide us with the means to connect theoretical predictions with actual experimental observations.

² To be fair, in the following sentence Graham says “you need to know how to calculate time and space complexity and about Turing completeness”. Apparently, NP-hardness, randomization, cryptography, and quantum computing are not essential to a hacker’s education.

The origin of much of science and medicine can be traced back to the ancient Babylonians. But perhaps their greatest contribution to humanity was the invention of the *place-value number system*. This is the idea that we can represent any number using a fixed number of digits, whereby the *position* of the digit is used to determine the corresponding value, as opposed to system such as Roman numerals, where every symbol has a fixed numerical value regardless of position. For example, the distance to the moon is 238,900 of our miles or 259,956 Roman miles. The latter quantity, expressed in standard Roman numerals is

MM
MM
MM

by a power of 10 (which is after all a simple shift) as well as multiply by a single-digit (which like addition, is done by multiplying each digit and propagating carries). Now suppose that x and y are two numbers of n decimal digits each. Adding two such numbers takes at least n single-digit additions (depending on how many times we need to use a “carry”), and so adding x to itself y times will take at least $n \cdot y$ single-digit additions. In contrast, the standard grade-school algorithm reduces this problem to taking n products of x with a single-digit (which require up to $2n$ single-digit operations each, depending on carries) and then adding all of those together (total of n additions, which, again depending on carries, would cost at most $2n^2$ single-digit operations) for a total of at most $4n^2$ single-digit operations. How much faster would $4n^2$ operations be than $n \cdot y$? And would this make any difference in a modern computer?

Let us consider the case of multiplying 64-bit or 20-digit numbers.⁴ That is, the task of multiplying two numbers x, y that are between 10^{19} and 10^{20} . Since in this case $n = 20$, the standard algorithm would use at most $4n^2 = 1600$ single-digit operations, while repeated addition would require at least $n \cdot y \geq 20 \cdot 10^{19}$ single-digit operations. To understand the difference, consider that a human being might do a single-digit operation in about 2 seconds, requiring just under an hour to complete the calculation of $x \times y$ using the grade-school algorithm. In contrast, even though it is more than a billion times faster, a modern PC that computes $x \times y$ using naïve iterated addition would require about $10^{20}/10^9 = 10^{11}$ seconds (which is more than three millenia!) to compute the same result.

⁴ This is a common size in several programming languages; for example the `long` data type in the Java programming language, and (depending on architecture) the `long` or `long long` types in C.

P It is important to distinguish between the *value* of a number, and the *length of its representation* (i.e., the number of digits it has). There is a big difference between the two: having 1,000,000,000 dollars is not the same as having 10 dollars! When talking about running time of algorithms, “less is more”, and so an algorithm that runs in time proportional to the *number of digits* of an input number (or even the number of digit squared) is much preferred to an algorithm that runs in time proportional to the *value* of the input number.

We see that computers have not made algorithms obsolete. On the contrary, the vast increase in our ability to measure, store, and communicate data has led to a much higher demand for developing better and more sophisticated algorithms that can allow us to make better decisions based on these data. We also see that to a large extent the notion of *algorithm* is independent of the actual computing device that will execute it. The digit-by-digit multiplication algorithm is

vastly better than iterated addition, regardless whether the technology we use to implement it is a silicon based chip, or a third grader with pen and paper.

Theoretical computer science is concerned with the *inherent* properties of algorithms and computation; namely, those properties that are *independent* of current technology. We ask some questions that were already pondered by the Babylonians, such as “what is the best way to multiply two numbers?”, but also questions that rely on cutting-edge science such as “could we use the effects of quantum entanglement to factor numbers faster?”.

In Computer Science parlance, a scheme such as the decimal (or sexadecimal) positional representation for numbers is known as a *data structure*, while the operations on this representations are known as *algorithms*. Data structures and algorithms have enabled amazing applications, but their importance goes beyond their practical utility. Structures from computer science, such as bits, strings, graphs, and even the notion of a program itself, as well as concepts such as universality and replication, have not just found (many) practical uses but contributed a new language and a new way to view the world.

0.1 EXTENDED EXAMPLE: A FASTER WAY TO MULTIPLY

Once you think of the standard digit-by-digit multiplication algorithm, it seems like obviously the “right” way to multiply numbers. Indeed, in 1960, the famous mathematician Andrey Kolmogorov organized a seminar at Moscow State University in which he conjectured that every algorithm for multiplying two n digit numbers would require a number of basic operations that is proportional to n^2 .⁵ Another way to say it, is that he conjectured that in any multiplication algorithm, doubling the number of digits would *quadruple* the number of basic operations required.

A young student named Anatoly Karatsuba was in the audience, and within a week he found an algorithm that requires only about $Cn^{1.6}$ operations for some constant C . Such a number becomes much smaller than n^2 as n grows.⁶ Amazingly, Karatsuba’s algorithm is based on a faster way to multiply *two-digit* numbers.

Suppose that $x, y \in [100] = \{0, \dots, 99\}$ are a pair of two-digit numbers. Let’s write \bar{x} for the “tens” digit of x , and \underline{x} for the “ones” digit, so that $x = 10\bar{x} + \underline{x}$, and write similarly $y = 10\bar{y} + \underline{y}$ for $\bar{y}, \underline{y} \in [10]$. The grade-school algorithm for multiplying x and y is illustrated in Fig. 1.

The grade-school algorithm works by transforming the task of multiplying a pair of two-digit number into *four* single-digit multipli-

⁵ That is, he conjectured that the number of operations would be at least some n^2/C operations for some constant C or, using “Big- O notation”, $\Omega(n^2)$ operations. See the *mathematical background* chapter for a precise definition of Big- O notation.

⁶ At the time of this writing, the **standard Python multiplication implementation** switches from the elementary school algorithm to Karatsuba’s algorithm when multiplying numbers larger than 1000 bits long.

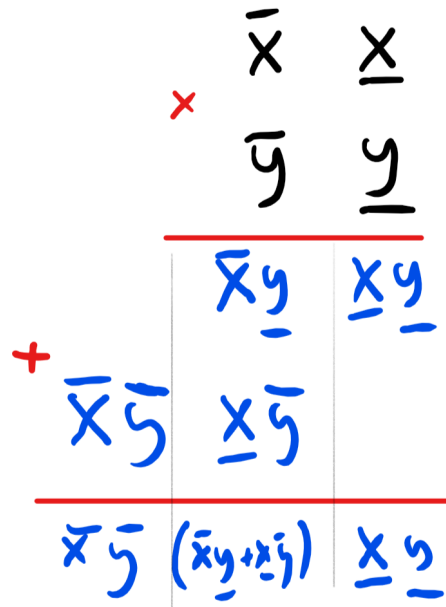


Figure 1: The grade-school multiplication algorithm illustrated for multiplying $x = 10\bar{x} + \underline{x}$ and $y = 10\bar{y} + \underline{y}$. It uses the formula $(10\bar{x} + \underline{x}) \times (10\bar{y} + \underline{y}) = 100\bar{x}\bar{y} + 10(\bar{x}\underline{y} + \underline{x}\bar{y}) + \underline{x}\underline{y}$.

cations via the formula

$$(10\bar{x} + \underline{x}) \times (10\bar{y} + \underline{y}) = 100\bar{x}\bar{y} + 10(\bar{x}\underline{y} + \underline{x}\bar{y}) + \underline{x}\underline{y} \quad (1)$$

Karatsuba’s algorithm is based on the observation that we can express this also as

$$(10\bar{x} + \underline{x}) \times (10\bar{y} + \underline{y}) = (100 - 10)\bar{x}\bar{y} + 10[(\bar{x} + \underline{x})(\bar{y} + \underline{y})] - (10 - 1)\underline{x}\underline{y} \quad (2)$$

which reduces multiplying the two-digit number x and y to computing the following three “simple” products: $\bar{x}\bar{y}$, $\underline{x}\underline{y}$ and $(\bar{x} + \underline{x})(\bar{y} + \underline{y})$.⁷

Of course if all we wanted to was to multiply two digit numbers, we wouldn’t really need any clever algorithms. It turns out that we can repeatedly apply the same idea, and use them to multiply 4-digit numbers, 8-digit numbers, 16-digit numbers, and so on and so forth. If we used the grade-school approach then our cost for doubling the number of digits would be to *quadruple* the number of multiplications, which for $n = 2^\ell$ digits would result in about $4^\ell = n^2$ operations. In contrast, in Karatsuba’s approach doubling the number of digits only *triples* the number of operations, which means that for $n = 2^\ell$ digits we require about $3^\ell = n^{\log_2 3} \sim n^{1.58}$ operations.

Specifically, we use a *recursive* strategy as follows:

⁷ The term $(\bar{x} + \underline{x})(\bar{y} + \underline{y})$ is not exactly a single-digit multiplication as $\bar{x} + \underline{x}$ and $\bar{y} + \underline{y}$ are numbers between 0 and 18 and not between 0 and 9. As we’ll see, it turns out that this does not make much of a difference, since when we use this algorithm recursively on n -digit numbers, this term will have at most $\lceil n/2 \rceil + 1$ digits, which is essentially half the number of digits as the original input.

$$\begin{array}{r}
 \bar{x} \quad \underline{x} \\
 \times \quad \underline{y} \quad \underline{y} \\
 \hline
 (\bar{x} + \underline{x})(\bar{y} + \underline{y}) \quad \underline{xy} \\
 + \quad \bar{x}\bar{y} \quad -\bar{x}\bar{y} - \underline{xy} \\
 \hline
 \bar{x}\bar{y} \quad (\bar{x}\underline{y} + \underline{x}\bar{y}) \quad \underline{xy}
 \end{array}$$

Figure 2: Karatsuba's multiplication algorithm illustrated for multiplying $x = 10\bar{x} + \underline{x}$ and $y = 10\bar{y} + \underline{y}$. We compute the three orange, green and purple products \underline{xy} , $\bar{x}\bar{y}$ and $(\bar{x} + \underline{x})(\bar{y} + \underline{y})$ and then add and subtract them to obtain the result.

Karatsuba Multiplication:

Input: nonnegative integers x, y each of at most n digits

Operation:

1. If $n \leq 2$ then return $x \cdot y$ (using a constant number of single-digit multiplications)
2. Otherwise, let $m = \lfloor n/2 \rfloor$, and write $x = 10^m \bar{x} + \underline{x}$ and $y = 10^m \bar{y} + \underline{y}$.⁸
2. Use recursion to compute $A = \bar{x}\bar{y}$, $B = \underline{xy}$ and $C = (\bar{x} + \underline{x})(\bar{y} + \underline{y})$. Note that all the numbers will have at most $m + 1$ digits.
3. Return $(10^n - 10^m) \cdot A + 10^m \cdot B + (1 - 10^m) \cdot C$

To understand why the output will be correct, first note that for $n > 2$, it will always hold that $m < n - 1$, and hence the recursive calls will always be for multiplying numbers with a smaller number of digits, and (since eventually we will get to single or double digit numbers) the algorithm will indeed terminate. Now, since $x = 10^m \bar{x} + \underline{x}$ and $y = 10^m \bar{y} + \underline{y}$,

$$x \times y = 10^n \bar{x} \cdot \bar{y} + 10^m (\bar{x}\bar{y} + \underline{xy}) + \underline{xy}. \quad (3)$$

⁸ Recall that for a number x , $\lfloor x \rfloor$ is obtained by "rounding down" x to the largest integer smaller or equal to x .

Rearranging the terms we see that

$$x \times y = 10^n \bar{x} \cdot \bar{y} + 10^m [(\bar{x} + \underline{x})(\bar{y} + \underline{y}) - \underline{xy} - \bar{x}\underline{y}] + \underline{xy}, \quad (4)$$

which equals $(10^n - 10^m) \cdot A + 10^m \cdot B + (1 - 10^m) \cdot C$, the value returned by the algorithm.

The key observation is that Eq. (4) reduces the task of computing the product of two n -digit numbers to computing *three* products of $\lceil n/2 \rceil$ -digit numbers. Specifically, we can compute $x \times y$ from the three products $\bar{x}\bar{y}$, \underline{xy} and $(\bar{x} + \underline{x})(\bar{y} + \underline{y})$, using a constant number (in fact eight) of additions, subtractions, and multiplications by 10^n or $10^{\lfloor n/2 \rfloor}$. (Multiplication by a power of ten can be done very efficiently as it corresponds to simply shifting the digits.) Intuitively this means that as the number of digits *doubles*, the cost of performing a multiplication via Karatsuba's algorithm *triples* instead of *quadruples*, as happens in the naive algorithm. This implies that multiplying numbers of $n = 2^\ell$ digits costs about $3^\ell = n^{\log_2 3} \sim n^{1.585}$ operations. In a [Exercise 0.3](#), you will formally show that the number of single-digit operations that Karatsuba's algorithm uses for multiplying n digit integers is at most $O(n^{\log_2 3})$ (see also [Fig. 2](#)).

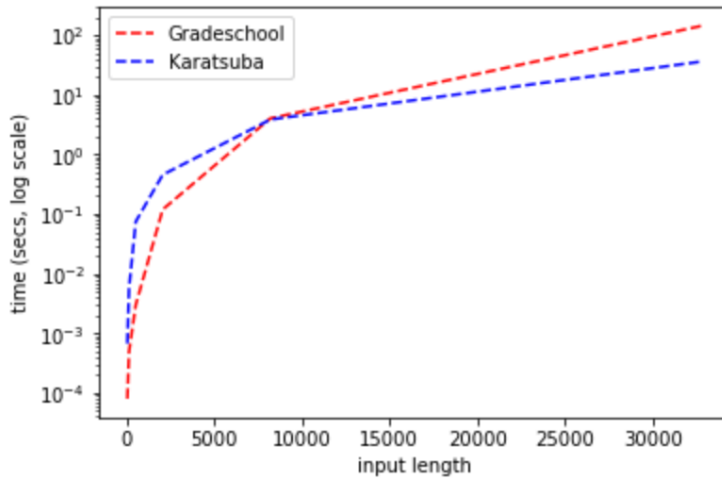


Figure 3: Running time of Karatsuba's algorithm vs. the grade-school algorithm. (Python implementation available [online](#).) Note the existence of a "cutoff" length, where for sufficiently large inputs Karatsuba becomes more efficient than the grade-school algorithm. The precise cutoff location varies by implementation and platform details, but will always occur eventually.

R **Ceilings, floors, and rounding** One of the benefits of using Big- O notation is that we can allow ourselves to be a little looser with issues such as rounding numbers etc.. For example, the natural way to describe Karatsuba's algorithm's running time is via




Figure 4: Karatsuba's algorithm reduces an n -bit multiplication to three $n/2$ -bit multiplications, which in turn are reduced to nine $n/4$ -bit multiplications and so on. We can represent the computational cost of all these multiplications in a 3-ary tree of depth $\log_2 n$, where at the root the extra cost is cn operations, at the first level the extra cost is $c(n/2)$ operations, and at each of the 3^i nodes of level i , the extra cost is $c(n/2^i)$.

The total cost is $cn \sum_{i=0}^{\log_2 n} (3/2)^i \leq 2cn^{\log_2 3}$ by the formula for summing a geometric series.

the following recursive equation

$$T(n) = 3T(n/2) + O(n) \quad (5)$$

but of course if n is not even then we cannot recursively invoke the algorithm on $n/2$ -digit integers. Rather, the true recursion is $T(n) = 3T(\lfloor n/2 \rfloor + 1) + O(n)$. However, this will not make much difference when we don't worry about constant factors, since it's not hard to show that $T(n + O(1)) \leq T(n) + o(T(n))$ for the functions we care about. Another way to show that this doesn't hurt us is to note that for every number n , we can find a number $n' \leq 2n$, such that n' is a power of two. Thus we can always "pad" the input by adding some input bits to make sure the number of digits is a power of two, in which case we will never run into these rounding issues. These kind of tricks work not just in the context of multiplication algorithms but in many other cases as well. Thus most of the time we can safely ignore these kinds of "rounding issues".

0.1.1 Beyond Karatsuba's algorithm

It turns out that the ideas of Karatsuba can be further extended to yield asymptotically faster multiplication algorithms, as was shown by Toom and Cook in the 1960s. But this was not the end of the line. In 1971, Schönhage and Strassen gave an even faster algorithm using the *Fast Fourier Transform*; their idea was to somehow treat integers as "signals" and do the multiplication more efficiently by moving to the

Fourier domain.⁹ The latest asymptotic improvement was given by Fürer in 2007 (though it only starts beating the Schönhage-Strassen algorithm for truly astronomical numbers). And yet, despite all this progress, we still don't know whether or not there is an $O(n)$ time algorithm for multiplying two n digit numbers!

⁹ The *Fourier transform* is a central tool in mathematics and engineering, used in a great number of applications. If you have not seen it yet, you will hopefully encounter it at some point in your studies.

R **Matrix Multiplication (advanced note)** (We will have several such “advanced” or “optional” notes and sections throughout this book. These may assume background that not every student has, and can be safely skipped over as none of the future parts will depend on them.)

It turns out that a similar idea as Karatsuba's can be used to speed up *matrix* multiplications as well. Matrices are a powerful way to represent linear equations and operations, widely used in a great many applications of scientific computing, graphics, machine learning, and many many more.

One of the basic operations one can do with two matrices is to *multiply* them. For example,

$$\text{if } x = \begin{pmatrix} x_{0,0} & x_{0,1} \\ x_{1,0} & x_{1,1} \end{pmatrix} \text{ and } y = \begin{pmatrix} y_{0,0} & y_{0,1} \\ y_{1,0} & y_{1,1} \end{pmatrix}$$

then the product of x and y is the matrix

$$\begin{pmatrix} x_{0,0}y_{0,0} + x_{0,1}y_{1,0} & x_{0,0}y_{0,1} + x_{0,1}y_{1,1} \\ x_{1,0}y_{0,0} + x_{1,1}y_{1,0} & x_{1,0}y_{0,1} + x_{1,1}y_{1,1} \end{pmatrix}. \text{ You}$$

can see that we can compute this matrix by *eight* products of numbers.

Now suppose that n is even and x and y are a pair of $n \times n$ matrices which we can think of as each composed of four $(n/2) \times (n/2)$ blocks $x_{0,0}, x_{0,1}, x_{1,0}, x_{1,1}$ and $y_{0,0}, y_{0,1}, y_{1,0}, y_{1,1}$. Then the formula for the matrix product of x and y can be expressed in the same way as above, just replacing products $x_{a,b}y_{c,d}$ with *matrix* products, and addition with matrix addition. This means that we can use the formula above to give an algorithm that *doubles* the dimension of the matrices at the expense of increasing the number of operation by a factor of 8, which for $n = 2^\ell$ will result in $8^\ell = n^3$ operations.

In 1969 Volker Strassen noted that we can compute the product of a pair of two-by-two matrices using only *seven* products of numbers by observing that each entry of the matrix xy can be computed by adding and subtracting the following seven terms:

$$\begin{aligned} t_1 &= (x_{0,0} + x_{1,1})(y_{0,0} + y_{1,1}), & t_2 &= (x_{0,0} + x_{1,1})y_{0,0}, \\ t_3 &= x_{0,0}(y_{0,1} - y_{1,1}), & t_4 &= x_{1,1}(y_{0,1} - y_{0,0}), \\ t_5 &= (x_{0,0} + x_{0,1})y_{1,1}, & t_6 &= (x_{1,0} - x_{0,0})(y_{0,0} + y_{0,1}), \\ t_7 &= (x_{0,1} - x_{1,1})(y_{1,0} + y_{1,1}). \end{aligned}$$

$$\text{that } xy = \begin{pmatrix} t_1 + t_4 - t_5 + t_7 & t_3 + t_5 \\ t_2 + t_4 & t_1 + t_3 - t_2 + t_6 \end{pmatrix}.$$

Using this observation, we can obtain an algorithm

such that doubling the dimension of the matrices results in increasing the number of operations by a factor of 7, which means that for $n = 2^\ell$ the cost is $7^\ell = n^{\log_2 7} \sim n^{2.807}$. A long sequence of work has since improved this algorithm, and the **current record** has running time about $O(n^{2.373})$. However, unlike the case of integer multiplication, at the moment we don't know of any algorithm for matrix multiplication that runs in time linear or even close to linear in the size of the input matrices (e.g., an $O(n^2 \text{polylog}(n))$ time algorithm). People have tried to use **group representations**, which can be thought of as generalizations of the Fourier transform, to obtain faster algorithms, but this effort **has not yet succeeded**.

0.2 ALGORITHMS BEYOND ARITHMETIC

The quest for better algorithms is by no means restricted to arithmetical tasks such as adding, multiplying or solving equations. Many *graph algorithms*, including algorithms for finding paths, matchings, spanning trees, cuts, and flows, have been discovered in the last several decades, and this is still an intensive area of research. (For example, the last few years saw many advances in algorithms for the *maximum flow* problem, borne out of surprising connections with electrical circuits and linear equation solvers.) These algorithms are being used not just for the “natural” applications of routing network traffic or GPS-based navigation, but also for applications as varied as drug discovery through searching for structures in gene-interaction graphs to computing risks from correlations in financial investments.

Google was founded based on the *PageRank* algorithm, which is an efficient algorithm to approximate the “principal eigenvector” of (a dampened version of) the adjacency matrix of web graph. The *Akamai* company was founded based on a new data structure, known as *consistent hashing*, for a hash table where buckets are stored at different servers. The *backpropagation algorithm*, which computes partial derivatives of a neural network in $O(n)$ instead of $O(n^2)$ time, underlies many of the recent phenomenal successes of learning deep neural networks. Algorithms for solving linear equations under sparsity constraints, a concept known as *compressed sensing*, have been used to drastically reduce the amount and quality of data needed to analyze MRI images. This is absolutely crucial for MRI imaging of cancer tumors in children, where previously doctors needed to use anesthesia to suspend breath during the MRI exam, sometimes with dire consequences.

Even for classical questions, studied through the ages, new dis-

coveries are still being made. For example, for the question of determining whether a given integer is prime or composite, which has been studied since the days of Pythagoras, efficient probabilistic algorithms were only discovered in the 1970s, while the first **deterministic polynomial-time algorithm** was only found in 2002. For the related problem of actually finding the factors of a composite number, new algorithms were found in the 1980s, and (as we'll see later in this course) discoveries in the 1990s raised the tantalizing prospect of obtaining faster algorithms through the use of quantum mechanical effects.

Despite all this progress, there are still many more questions than answers in the world of algorithms. For almost all natural problems, we do not know whether the current algorithm is the “best”, or whether a significantly better one is still waiting to be discovered. As we already saw, even for the classical problem of multiplying numbers we have not yet answered the age-old question of **“is multiplication harder than addition?”** .

But at least we now know the right way to *ask* it.

0.3 ON THE IMPORTANCE OF NEGATIVE RESULTS.

Finding better multiplication algorithms is undoubtedly a worthwhile endeavor. But why is it important to prove that such algorithms *don't* exist? What useful applications could possibly arise from an impossibility result?

One motivation is pure intellectual curiosity. After all, this is a question even Archimedes could have been excited about. Another reason to study impossibility results is that they correspond to the fundamental limits of our world. In other words, they are *laws of nature*. In physics, the impossibility of building a *perpetual motion machine* corresponds to the *law of conservation of energy*. The impossibility of building a heat engine beating Carnot's bound corresponds to the second law of thermodynamics, while the impossibility of faster-than-light information transmission is a cornerstone of special relativity.

In mathematics, while we all learned the solution for quadratic equations in high school, the impossibility of generalizing this to equations of degree five or more gave birth to *group theory*. Another example of an impossibility result comes from geometry. For two millennia, mathematicians tried to show that Euclid's fifth axiom or “postulate” could be derived from the first four. (This fifth postulate was known as the “parallel postulate”, and roughly speaking it states that every line has a unique parallel line of each distance.) It was shown to be impossible using constructions of so called “non-Euclidean geometries”, which turn out to be crucial for the theory of general relativity.

R It is fine if you have not yet encountered many of the above examples. I hope however that they spark your curiosity!

In an analogous way, impossibility results for computation correspond to “computational laws of nature” that tell us about the fundamental limits of any information processing apparatus, whether based on silicon, neurons, or quantum particles.¹⁰ Moreover, computer scientists have recently been finding creative approaches to *apply* computational limitations to achieve certain useful tasks. For example, much of modern Internet traffic is encrypted using the RSA encryption scheme, which relies on its security on the (conjectured) impossibility of efficiently factoring large integers. More recently, the **Bitcoin** system uses a digital analog of the “gold standard” where, instead of using a precious metal, new currency is obtained by “mining” solutions for computationally difficult problems.

¹⁰ Indeed, some **exciting recent research** is focused on trying to use computational complexity to shed light on fundamental questions in physics such as understanding black holes and reconciling general relativity with quantum mechanics.

✓ Lecture Recap

- The history of algorithms goes back thousands of years; they have been essential much of human progress and these days form the basis of multi-billion dollar industries, as well as life-saving technologies.
- There is often more than one algorithm to achieve the same computational task. Finding a faster algorithm can often make a much bigger difference than improving computing hardware.
- Better algorithms and data structures don’t just speed up calculations, but can yield new qualitative insights.
- One question we will study is to find out what is the *most efficient* algorithm for a given problem.
- To show that an algorithm is the most efficient one for a given problem, we need to be able to *prove* that it is *impossible* to solve the problem using a smaller amount of computational resources.

0.4 ROADMAP TO THE REST OF THIS COURSE

Often, when we try to solve a computational problem, whether it is solving a system of linear equations, finding the top eigenvector of a matrix, or trying to rank Internet search results, it is enough to use the “I know it when I see it” standard for describing algorithms. As long as we find some way to solve the problem, we are happy and don’t care so much about formal descriptions of the algorithm. But when

we want to answer a question such as “does there *exist* an algorithm to solve the problem P ?” we need to be much more precise.

In particular, we will need to (1) define exactly what it means to solve P , and (2) define exactly what an algorithm is. Even (1) can sometimes be non-trivial but (2) is particularly challenging; it is not at all clear how (and even whether) we can encompass all potential ways to design algorithms. We will consider several simple *models of computation*, and argue that, despite their simplicity, they do capture all “reasonable” approaches to achieve computing, including all those that are currently used in modern computing devices.

Once we have these formal models of computation, we can try to obtain *impossibility results* for computational tasks, showing that some problems *can not be solved* (or perhaps can not be solved within the resources of our universe). Archimedes once said that given a fulcrum and a long enough lever, he could move the world. We will see how *reductions* allow us to leverage one hardness result into a slew of a great many others, illuminating the boundaries between the computable and uncomputable (or tractable and intractable) problems.

Later in this course we will go back to examining our models of computation, and see how resources such as randomness or quantum entanglement could potentially change the power of our model. In the context of probabilistic algorithms, we will see a glimpse of how randomness has become an indispensable tool for understanding computation, information, and communication. We will also see how computational difficulty can be an asset rather than a hindrance, and be used for the “derandomization” of probabilistic algorithms. The same ideas also show up in *cryptography*, which has undergone not just a technological but also an intellectual revolution in the last few decades, much of it building on the foundations that we explore in this course.

Theoretical Computer Science is a vast topic, branching out and touching upon many scientific and engineering disciplines. This course only provides a very partial (and biased) sample of this area. More than anything, I hope I will manage to “infect” you with at least some of my love for this field, which is inspired and enriched by the connection to practice, but which I find to be deep and beautiful regardless of applications.

0.4.1 Dependencies between chapters

This book is divided into the following parts:

- **Preliminaries:** Introduction, mathematical background, and representing objects as strings.
- **Part I: Finite computation:** Boolean circuits / straightline pro-

grams. Universal gatesets, counting lower bound, representing programs as string and universality.

- **Part II: Uniform computation:** Turing machines / programs with loops. Equivalence of models (including RAM machines and λ calculus), universality, uncomputability, Gödel's incompleteness theorem, restricted models (regular and context free languages).
- **Part III: Efficient computation:** Definition of running time, time hierarchy theorem, P and NP , NP completeness, space bounded computation.
- **Part IV: Randomized computation:** Probability, randomized algorithms, BPP , amplification, $BPP \subseteq P_{poly}$, pseudorandom generators and derandomization.
- **Part V: Advanced topics:** Cryptography, proofs and algorithms (interactive and zero knowledge proofs, Curry-Howard correspondence), quantum computing.

The book proceeds in linear order, with each chapter building on the previous one, with the following exceptions:

- All chapters in [Part V](#) (Advanced topics) are independent of one another, and you can choose which one of them to read.
- [Chapter 10](#) (Gödel's incompleteness theorem), [Chapter 9](#) (Restricted computational models), and [Chapter 16](#) (Space bounded computation), are not used in following chapters. Hence you can choose to skip them.

A course based on this book can use all of Parts I, II, and III (possibly skipping over some or all of [Chapter 10](#), [Chapter 9](#) or [Chapter 16](#)), and then either cover all or some of Part IV, and add a "sprinkling" of advanced topics from Part V based on student or instructor interest.

0.5 EXERCISES

R Disclaimer Most of the exercises have been written in the summer of 2018 and haven't yet been fully debugged. While I would prefer people do not post online solutions to the exercises, I would greatly appreciate if you let me know of any bugs. You can do so by posting a [GitHub issue](#) about the exercise, and optionally complement this with an email to me with more details about the attempted solution.

Exercise 0.1 Rank the significance of the following inventions in speeding up multiplication of large (that is 100-digit or more) numbers. That is, use “back of the envelope” estimates to order them in terms of the speedup factor they offered over the previous state of affairs.

1. Discovery of the grade-school digit by digit algorithm (improving upon repeated addition)
2. Discovery of Karatsuba’s algorithm (improving upon the digit by digit algorithm)
3. Invention of modern electronic computers (improving upon calculations with pen and paper).

Exercise 0.2 The 1977 Apple II personal computer had a processor speed of 1.023 Mhz or about 10^6 operations per seconds. At the time of this writing the world’s fastest supercomputer performs 93 “petaflops” (10^{15} floating point operations per second) or about 10^{18} basic steps per second. For each one of the following running times (as a function of the input length n), compute for both computers how large an input they could handle in a week of computation, if they run an algorithm that has this running time:

1. n operations.
2. n^2 operations.
3. $n \log n$ operations.
4. 2^n operations.
5. $n!$ operations.

Exercise 0.3 — Analysis of Karatsuba’s Algorithm. 1. Suppose that T_1, T_2, T_3, \dots is a sequence of numbers such that $T_2 \leq 10$ and for every n , $T_n \leq 3T_{\lceil n/2 \rceil} + Cn$. Prove that $T_n \leq 10Cn^{\log_2 3}$ for every n .¹¹

2. Prove that the number of single-digit operations that Karatsuba’s algorithm takes to multiply two n digit numbers is at most $1000n^{\log_2 3}$.

¹¹ **Hint:** Use a proof by induction - suppose that this is true for all n ’s from 1 to m , prove that this is true also for $m + 1$.

Exercise 0.4 Implement in the programming language of your choice functions `Gradeschool_multiply(x, y)` and `Karatsuba_multiply(x, y)` that take two arrays of digits x and y and

return an array representing the product of x and y (where x is identified with the number $x[0]+10*x[1]+100*x[2]+\dots$ etc..) using the grade-school algorithm and the Karatsuba algorithm respectively. At what number of digits does the Karatsuba algorithm beat the grade-school one? ■

Exercise 0.5 — Matrix Multiplication (optional, advanced). In this exercise, we show that if for some $\omega > 2$, we can write the product of two $k \times k$ real-valued matrices A, B using at most k^ω multiplications, then we can multiply two $n \times n$ matrices in roughly n^ω time for every large enough n .

To make this precise, we need to make some notation that is unfortunately somewhat cumbersome. Assume that there is some $k \in \mathbb{N}$ and $m \leq k^\omega$ such that for every $k \times k$ matrices A, B, C such that $C = AB$, we can write for every $i, j \in [k]$:

$$C_{i,j} = \sum_{\ell=0}^m \alpha_{i,j}^\ell f_\ell(A) g_\ell(B) \quad (6)$$

for some linear functions $f_0, \dots, f_{m-1}, g_0, \dots, g_{m-1} : \mathbb{R}^{n^2} \rightarrow \mathbb{R}$ and coefficients $\{\alpha_{i,j}^\ell\}_{i,j \in [k], \ell \in [m]}$. Prove that under this assumption for every $\epsilon > 0$, if n is sufficiently large, then there is an algorithm that computes the product of two $n \times n$ matrices using at most $O(n^{\omega+\epsilon})$ arithmetic operations.¹² ■

¹² Hint: Start by showing this for the case that $n = k^t$ for some natural number t , in which case you can do so recursively by breaking the matrices into $k \times k$ blocks.

0.6 BIBLIOGRAPHICAL NOTES

For an overview of what we'll see in this course, you could do far worse than read [Bernard Chazelle's wonderful essay on the Algorithm as an Idiom of modern science](#).

0.7 FURTHER EXPLORATIONS

Some topics related to this chapter that might be accessible to advanced students include:

- The *Fourier transform*, the *Fast Fourier transform* algorithm and how to use it multiply polynomials and integers. [This lecture of Jeff Erickson](#) (taken from his [collection of notes](#)) is a very good starting point. See also this [MIT lecture](#) and this [popular article](#).
- Fast matrix multiplication algorithms, and the approach of obtaining exponent two via group representations.
- The proofs of some of the classical impossibility results in mathematics we mentioned, including the impossibility of proving Euclid's fifth postulate from the other four, impossibility of trisecting an angle with a straightedge and compass and the impossibility

of solving a quintic equation via radicals. A geometric proof of the impossibility of angle trisection (one of the three [geometric problems of antiquity](#), going back to the ancient greeks) is given in [this blog post of Tao](#). [This book of Mario Livio](#) covers some of the background and ideas behind these impossibility results.

