
Computational Complexity: A Modern Approach

Draft of a book: Dated January 2007
Comments welcome!

Sanjeev Arora and Boaz Barak
Princeton University
complexitybook@gmail.com

Not to be reproduced or distributed without the authors' permission

This is an Internet draft. Some chapters are more finished than others. References and attributions are very preliminary and we apologize in advance for any omissions (but hope you will nevertheless point them out to us).

**Please send us bugs, typos, missing references or general comments to
complexitybook@gmail.com — Thank You!!**

Chapter 1

The computational model —and why it doesn't matter

“The idea behind digital computers may be explained by saying that these machines are intended to carry out any operations which could be done by a human computer. The human computer is supposed to be following fixed rules; he has no authority to deviate from them in any detail. We may suppose that these rules are supplied in a book, which is altered whenever he is put on to a new job. He has also an unlimited supply of paper on which he does his calculations.”

Alan Turing, 1950

“[Turing] has for the first time succeeded in giving an absolute definition of an interesting epistemological notion, i.e., one not depending on the formalism chosen.”

Kurt Gödel, 1946

The problem of mathematically modeling computation may at first seem insurmountable: throughout history people have been solving computational tasks using such a wide variety of methods, ranging from intuition and “eureka” moments, to mechanical devices such as abacus or sliderulers, to modern computers. Besides that, other organisms and systems in nature are also faced with and solve computational tasks everyday using a bewildering array of tools. How can you find a simple mathematical model that captures all of these ways to compute? The problem is even more exacerbated since in this book we are interested in issues of *computational efficiency*. Here, at first glance, it seems that we have to be very careful about our choice of a computational model, since even a kid knows that whether or not a new video game program is “efficiently computable” depends upon his computer’s hardware.

Surprisingly enough it turns out there there is a simple mathematical model that suffices for studying many questions about computation and its efficiency —the *Turing machine*. It suffices to restrict attention to this single model since it seems able to *simulate* all physically realizable computational methods with very little loss of efficiency. Thus the set of “efficiently solvable” computational tasks is at least as large for the Turing Machine as for any other method of computation. (One possible exception is the quantum computer model described in Chapter 10, but we do not currently know if it is physically realizable.)

In this chapter we formally define Turing machines and survey some of their basic properties. Section 1.1, presents the formal definition of the Turing machine. Section 1.3 presents one of the most profound facts about Turing machines (and algorithms in general): they can be represented as *strings* (i.e., finite sequence of symbols) and then be used as *inputs* for other algorithms or Turing machines. One surprising application of this fact appears in Section 1.4: the existence

of functions that are *uncomputable* by any algorithm, regardless of its running time. However, complexity theory is interested not so much in finding out what can be computed in principle, but rather in finding out what can be computed *efficiently*. Towards this end Section 1.5 contains one of the most important definitions in this book: the definition of complexity class \mathbf{P} , which aims to capture mathematically the set of all decision problems that can be efficiently solved. Section 1.5 also contains some discussion on whether or not the class \mathbf{P} truly captures the informal notion of “efficient computation”. The section also points out how throughout the book the definition of the Turing Machine and the class \mathbf{P} will be a starting point for definitions of many other models, including nondeterministic, probabilistic and quantum Turing machines, Boolean circuits, parallel computers, decision trees, and communication games. Some of these models are introduced to study arguably realizable modes of physical computation, while others are mainly used to gain insights on Turing machines.

1.1 Modeling computation

For thousands of years, the term “computation” was understood to mean application of mechanical rules to manipulate numbers, where the person/machine doing the manipulation is allowed a *scratch pad* on which to write the intermediate results. The Turing Machine is a concrete embodiment of this intuitive notion. It can be also viewed as the equivalent of any modern programming language — albeit one with no built-in prohibition about memory size.¹ In fact, this intuitive understanding of computation will suffice for most of the book and most readers can skip many details of the model on a first reading, returning to them later as needed.

We first recall the informal description of computation (see Turing’s quote at the start of the chapter). Let f be a function that takes a string of bits (i.e., a member of the set $\{0, 1\}^*$) and outputs, say, either 0 or 1. An *algorithm* for computing f is a set of mechanical rules, such that by following them we can compute $f(x)$ given any input $x \in \{0, 1\}^*$. The set of rules being followed is fixed (i.e., the same rules must work for all possible inputs) though each rule in this set may be applied arbitrarily many times. Each rule involves one or more of the following “elementary” operations:

1. Read a bit of the input.
2. Read a bit (or possibly a symbol from a slightly larger alphabet, say a digit in the set $\{0, \dots, 9\}$) from the “scratch pad” or working space we allow the algorithm to use.

Based on the values read,

3. Write a bit/symbol to the scratch pad.
4. Either stop and output 0 or 1, or choose a new rule from the set that will be applied next.

Finally, the *running time* is the number of these basic operations performed.

Below, we formalize all of these notions.

1.1.1 The Turing Machine

The *k-tape Turing machine* (TM) concretely realizes the above informal notion in the following way (see Figure 1.1):

Scratch Pad: The scratch pad consists of k tapes. A *tape* is an infinite one-directional line of cells, each of which can hold a symbol from a finite set Γ called the *alphabet* of the machine. Each

¹Though the assumption of a potentially infinite memory may seem unrealistic at first, in the complexity setting it is of no consequence since we will restrict the machine to use a finite amount of tape cells for any given input (the number allowed will depend upon the input size).

tape is equipped with a *tape head* that can potentially read or write symbols to the tape one cell at a time. The machine's computation is divided into discrete time steps, and the head can move left or right one cell in each step.

The first tape of the machine is designated as the *input* tape. The machine's head can only read symbols from that tape, not write them—a so-called read-only head. The $k - 1$ read-write tapes are called *work tapes* and the last one of them is designated as the *output tape* of the machine, on which it writes its final answer before halting its computation.

There also are variants of Turing machines with *random access memory*,² but it turns out that their computational powers are equivalent to standard Turing machines (see Exercise 1.13).

Finite set of operations/rules: The machine has a finite set of *states*, denoted Q . The machine contains a “register” that can hold a single element of Q ; this is the “state” of the machine at that instant. This state determines its action at the next computational step, which consists of the following: **(1)** read the symbols in the cells directly under the k heads **(2)** for the $k - 1$ read/write tapes replace each symbol with a new symbol (it has the option of not changing the tape by writing down the old symbol again), **(3)** change its register to contain another state from the finite set Q (it has the option not to change its state by choosing the old state again) and **(4)** move each head one cell to the left or to the right (or stay in place).

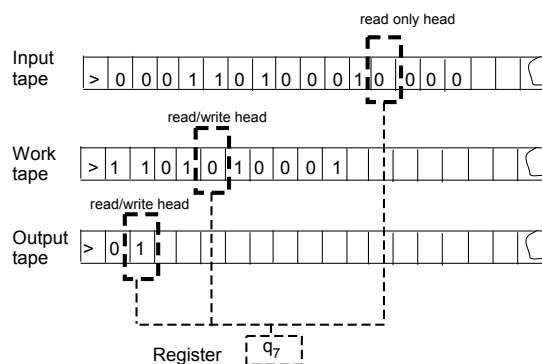


Figure 1.1: A snapshot of the execution of a 3-tape Turing machine M with an input tape, a work tape, and an output tape.

One can think of the Turing machine as a simplified modern computer, with the machine's tape corresponding to a computer's memory, and the transition function and register corresponding to the computer's central processing unit (CPU). However, it's best to think of Turing machines as simply a formal way to describe algorithms. Even though algorithms are often best described by plain English text, it is sometimes useful to express them by such a formalism in order to argue about them mathematically. (Similarly, one needs to express an algorithm in a programming language in order to execute it on a computer.)

Formal definition. Formally, a TM M is described by a tuple (Γ, Q, δ) containing:

- A finite set Γ of the symbols that M 's tapes can contain. We assume that Γ contains a designated “blank” symbol, denoted \square , a designated “start” symbol, denoted \triangleright and the numbers 0 and 1. We call Γ the *alphabet* of M .
- A finite set Q of possible states M 's register can be in. We assume that Q contains a designated start state, denoted q_{start} and a designated halting state, denoted q_{halt} .

²*Random access* denotes the ability to access the i^{th} symbol of the memory within a single step, without having to move a head all the way to the i^{th} location. The name “random access” is somewhat unfortunate since this concept involves no notion of randomness—perhaps “indexed access” would have been better. However, “random access” is widely used and so we follow this convention in this book.

- A function $\delta: Q \times \Gamma^k \rightarrow Q \times \Gamma^{k-1} \times \{L, S, R\}^k$, where $k \geq 2$, describing the rules M use in performing each step. This function is called the *transition function* of M (see Figure 1.2.)

IF			THEN			
input symbol read	work/output tape symbol read	current state	move input head	new work/output tape symbol	move work/output tape	new state
⋮	⋮	⋮	⋮	⋮	⋮	⋮
a	b	q	Right →	b'	← Left	q'
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Figure 1.2: The transition function of a two tape TM (i.e., a TM with one input tape and one work/output tape).

If the machine is in state $q \in Q$ and $(\sigma_1, \sigma_2, \dots, \sigma_k)$ are the symbols currently being read in the k tapes, and $\delta(q, (\sigma_1, \dots, \sigma_k)) = (q', (\sigma'_2, \dots, \sigma'_k), z)$ where $z \in \{L, S, R\}^k$, then at the next step the σ symbols in the last $k - 1$ tapes will be replaced by the σ' symbols, the machine will be in state q' , and the k heads will move Left, Right or Stay in place, as given by z . (If the machine tries to move left from the leftmost position of a tape then it will stay in place.)

All tapes except for the input are initialized in their first location to the *start* symbol \triangleright and in all other locations to the *blank* symbol \square . The input tape contains initially the start symbol \triangleright , a finite non-blank string x (“the input”), and the rest of its cells are initialized with the blank symbol \square . All heads start at the left ends of the tapes and the machine is in the special starting state q_{start} . This is called the *start configuration* of M on input x . Each step of the computation is performed by applying the function δ as described above. The special halting state q_{halt} has the property that once the machine is in q_{halt} , the transition function δ does not allow it to further modify the tape or change states. Clearly, if the machine enters q_{halt} then it has *halted*. In complexity theory we are typically only interested in machines that halt for every input in a finite number of steps.

EXAMPLE 1.1

Let PAL be the Boolean function defined as follows: for every $x \in \{0, 1\}^*$, PAL(x) is equal to 1 if x is a *palindrome* and equal to 0 otherwise. That is, PAL(x) = 1 if and only if x reads the same from left to right as from right to left (i.e., $x_1x_2\dots x_n = x_nx_{n-1}\dots x_1$). We now show a TM M that computes PAL within less than $3n$ steps.

Our TM M will use 3 tapes (input, work and output) and the alphabet $\{\triangleright, \square, 0, 1\}$. It operates as follows:

1. Copy the input to the read/write work tape.
2. Move the input-tape head to the beginning of the input.
3. Move the input-tape head to the right while moving the work-tape head to the left. If at any moment the machine observes two different values, it halts and output 0.
4. Halt and output 1.

We now describe the machine more formally: The TM M uses 5 states denoted by $\{q_{\text{start}}, q_{\text{copy}}, q_{\text{left}}, q_{\text{test}}, q_{\text{halt}}\}$. Its transition function is defined as follows:

1. On state q_{start} , move the input-tape head to the right, and move the work-tape head to the right while writing the start symbol \triangleright ; change the state to q_{copy} . (Unless we mention this explicitly, the function does not change the output tape's contents or head position.)
 2. On state q_{copy} :
 - If the symbol read from the input tape is not the blank symbol \square then move both the input-tape and work-tape heads to the right, writing the symbol from the input-tape on the work-tape; stay in the state q_{copy} .
 - If the symbol read from the input tape is the blank symbol \square , then move the input-tape head to the left, while keeping the work-tape head in the same place (and not writing anything); change the state to q_{left} .
 3. On state q_{left} :
 - If the symbol read from the input tape is not the start symbol \triangleright then move the input-head to the left, keeping the work-tape head in the same place (and not writing anything); stay in the state q_{left} .
 - If the symbol read from the input tape is the start symbol \triangleright then move the input-tape to the right and the work-tape head to the left (not writing anything); change to the state q_{test} .
 4. On state q_{test} :
 - If the symbol read from the input-tape is the blank symbol \square and the symbol read from the work-tape is the start symbol \triangleright then write 1 on the output tape and change state to q_{halt} .
 - Otherwise, if the symbols read from the input tape and the work tape are not the same then write 0 on the output tape and change state to q_{halt} .
 - Otherwise, if the symbols read from the input tape and the work tape are the same, then move the input-tape head to the right and the work-tape head to the left; stay in the state q_{test} .
-

Clearly, fully specifying a Turing machine is tedious and not always informative. While it is useful to work out one or two examples for yourself (see Exercise 1.4), in the rest of this book we avoid such overly detailed descriptions and specify TM's in a high level fashion. For readers who know how to write computer programs, Example 1.2 below should convince them that they know (in principle at least) how to design a Turing machine for any computational task for which they know how to write computer programs.

1.1.2 The expressive power of Turing machines.

At first sight, it may be unclear that Turing machines do indeed encapsulate our intuitive notion of computation. It may be useful to work through some simple examples, such as expressing the standard algorithms for addition and multiplication in terms of Turing machines computing the corresponding functions (see Exercise 1.4). Having done that, you may be ready for the next example; it outlines how you can translate a program in your favorite programming language into a Turing machine. (The reverse direction also holds: most programming languages can simulate a Turing machine.)

EXAMPLE 1.2

[Simulating a general programming language using Turing machines]

(This example assumes some background in computing.)

We give a hand-wavy proof that for any program written in any of the familiar programming languages such as C or Java, there is an equivalent Turing machine. First, recall that programs in these programming languages can be translated (the technical term is *compiled*) into an equivalent *machine language* program. This is a sequence of simple instructions to read from memory into one of a finite number of registers, write a register's contents to memory, perform basic arithmetic operations, such as adding two registers, and control instructions that perform actions conditioned on, say, whether a certain register is equal to zero.

All these operations can be easily simulated by a Turing machine. The memory and registers can be implemented using the machine's tapes, while the instructions can be encoded by the machine's transition function. For example, it's not hard to show TM's that add or multiply two numbers, or a two-tape TM that, if its first tape contains a number i in binary representation, can read or modify the i^{th} location of its second tape (e.g., see Exercise 1.12).

Exercise 1.14 asks you to give a more rigorous proof of such a simulation for a simple tailor-made programming language.

1.2 Efficiency and running time

Now we formalize the notion of running time. As every non-trivial computational task requires at least reading the entire input, we count the number of basic steps *as a function of the input length*.

DEFINITION 1.3 (COMPUTING A FUNCTION AND RUNNING TIME)

Let $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and let $T : \mathbb{N} \rightarrow \mathbb{N}$ be some functions, and let M be a Turing machine. We say that M *computes* f if for every $x \in \{0, 1\}^*$, whenever M is initialized to the start configuration on input x , then it halts with $f(x)$ written on its output tape. We say M *computes* f *in* $T(n)$ -*time*³ if its computation every x requires at most $T(|x|)$ steps.

EXAMPLE 1.4

It is easily checked that the Turing machine for palindrome recognition in Example 1.1 runs in $O(n)$ time.

Time-constructible functions. A function $T : \mathbb{N} \rightarrow \mathbb{N}$ is *time constructible* if $T(n) \geq n$ and there is a TM M that computes the function $x \mapsto \lfloor T(|x|) \rfloor$ in time $T(n)$. (As usual, $\lfloor T(|x|) \rfloor$ denotes the binary representation of the number $T(|x|)$.) Examples for time-constructible functions are n , $n \log n$, n^2 , 2^n . Almost all functions encountered in this book will be time-constructible and, to avoid annoying anomalies, we will restrict our attention to time bounds of this form. (The restriction $T(n) \geq n$ is to allow the algorithm time to read its input.)

³Formally we should write “ T -time” instead of “ $T(n)$ -time”, but we follow the convention of writing $T(n)$ to emphasize that T is applied to the input length.

1.2.1 Robustness of our definition

Most of the specific details of our definition of Turing machines are quite arbitrary. It is a simple exercise to see that most changes to the definition do not yield a substantially different model, since our model can simulate any of these new models. In context of computational complexity, however, we have to verify not only that one model can simulate another, but that it can do so efficiently. Now we state a few results of this type.

The following simple claims show that restricting the alphabet Γ to be $\{0, 1, \square, \triangleright\}$, restricting the machine to have a single work tape, or allowing the tapes to be infinite in both directions will not have a significant effect on the time to compute functions: (Below we provide only proof sketches for these claims— completing these sketches into full proofs is a very good way to gain intuition on Turing machines, see Exercises 1.5, 1.6 and 1.7.)

CLAIM 1.5

For every $f : \{0, 1\}^* \rightarrow \{0, 1\}$ and time-constructible $T : \mathbb{N} \rightarrow \mathbb{N}$, if f is computable in time $T(n)$ by a TM M using alphabet Γ then it is computable in time $4 \log |\Gamma| T(n)$ by a TM \tilde{M} using the alphabet $\{0, 1, \square, \triangleright\}$.




Figure 1.3: We can simulate a machine M using the alphabet $\{\triangleright, \square, a, b, \dots, z\}$ by a machine M' using $\{\triangleright, \square, 0, 1\}$ via encoding every tape cell of M using 5 cells of M' .

PROOF SKETCH: Let M be a TM with alphabet Γ , k tapes, and state set Q that computes the function f in $T(n)$ time. We describe an equivalent TM \tilde{M} computing f with alphabet $\{0, 1, \square, \triangleright\}$, k tapes and a set Q' of states. The idea behind the transformation is simple: one can encode any member of Γ using $\log |\Gamma|$ bits.⁴ Thus, each of \tilde{M} 's work tapes will simply encode one of M 's tapes: for every cell in M 's tape we will have $\log |\Gamma|$ cells in the corresponding tape of \tilde{M} (see Figure 1.3).

To simulate one step of M , the machine \tilde{M} will: **(1)** use $\log |\Gamma|$ steps to read from each tape the $\log |\Gamma|$ bits encoding a symbol of Γ **(2)** use its state register to store the symbols read, **(3)** use M 's transition function to compute the symbols M writes and M 's new state given this information, **(3)** store this information in its state register, and **(4)** use $\log |\Gamma|$ steps to write the encodings of these symbols on its tapes.

One can verify that this can be carried out if \tilde{M} has access to registers that can store M 's state, k symbols in Γ and a counter from 1 to k . Thus, there is such a machine \tilde{M} utilizing no more than $c|Q||\Gamma|^k k$ states for some absolute constant c . (In general, we can always simulate several registers using one register with a larger state space. For example, we can simulate three registers taking values in the sets A, B and C respectively with one register taking a value in the set $A \times B \times C$ which is of size $|A||B||C|$.)

It is not hard to see that for every input $x \in \{0, 1\}^n$, if on input x the TM M outputs $f(x)$ within $T(n)$ steps, then \tilde{M} will output the same value within less than $4 \log |\Gamma| T(n)$ steps. ■

Now we consider the effect of restricting the machine to use a single tape (the standard computational model in many undergraduate texts)—one read/write tape that serves as input, work and output tape. We show that going from multiple tapes to a single tape can at most square the running time. (Furthermore, this quadratic increase actually happens for some languages, including palindrome recognition considered in Example 1.1; see chapter notes.)

⁴Recall our conventions that \log is taken to base 2, and non-integer numbers are rounded up when necessary.

CLAIM 1.6

For every $f : \{0, 1\}^* \rightarrow \{0, 1\}$ and time-constructible $T : \mathbb{N} \rightarrow \mathbb{N}$, if f is computable in time $T(n)$ by a TM M using k tapes then it is computable in time $5kT(n)^2$ by a TM \tilde{M} using only a single work tape (plus additional input and output tapes).

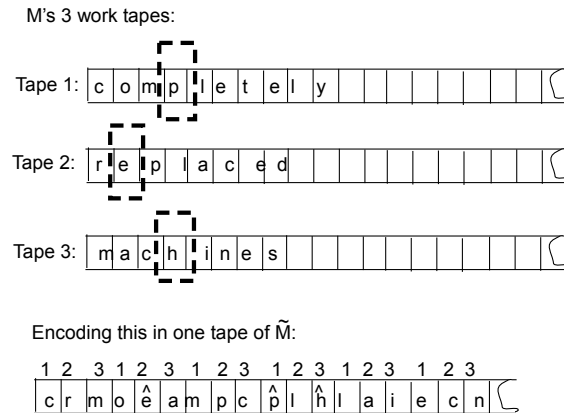


Figure 1.4: Simulating a machine M with 3 work tapes using a machine \tilde{M} with a single work tape (in addition to the input and output tapes).

PROOF SKETCH: Again the idea is simple: the TM \tilde{M} encodes the k tapes of M on a single tape by using locations $1, k + 1, 2k + 1, \dots$ to encode the first tape, locations $2, k + 2, 2k + 2, \dots$ to encode the second tape etc.. (see Figure 1.4). For every symbol a in M 's alphabet, \tilde{M} will contain both the symbol a and the symbol \hat{a} . In the encoding of each tape, exactly one symbol will be of the “ $\hat{\cdot}$ ” type, indicating that the corresponding head of M is positioned in that location (see figure). \tilde{M} uses the input and output tape in the same way M does. To simulate one step of M , the machine \tilde{M} makes two sweeps of its work tape: first it sweeps the tape in the left-to-right direction and records to its register the k symbols that are marked by “ $\hat{\cdot}$ ”. Then \tilde{M} uses M 's transition function to determine the new state, symbols, and head movements and sweeps the tape back in the right-to-left direction to update the encoding accordingly. Clearly, \tilde{M} will have the same output as M . Also, since on n -length inputs M never reaches more than location $T(n)$ of any of its tapes, \tilde{M} will never need to reach more than location $kT(n)$ of its work tape, meaning that for each of the at most $T(n)$ steps of M , \tilde{M} performs at most $5kT(n)$ work (sweeping back and forth requires about $2T(n)$ steps, and some additional steps may be needed for updating head movement and book keeping). ■

REMARK 1.7 (OBLIVIOUS TURING MACHINES.)

With a bit of care, one can ensure that the proof of Claim 1.6 yields a TM \tilde{M} with the following property: its head movements do not depend on the input but only on the input length. That is, every input $x \in \{0, 1\}^*$ and $i \in \mathbb{N}$, the location of each of M 's heads at the i^{th} step of execution on input x is only a function of $|x|$ and i . A machine with this property is called *oblivious* and the fact that every TM can be simulated by an oblivious TM will be useful for us later on (see Exercises 1.8, 1.9 and the proof of Theorem 2.10).

CLAIM 1.8

Define a bidirectional TM to be a TM whose tapes are infinite in both directions. For every $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ and time constructible $T : \mathbb{N} \rightarrow \mathbb{N}$, if f is computable in time $T(n)$ by a bidirectional TM M then it is computable in time $4T(n)$ by a standard (unidirectional) TM \tilde{M} .

PROOF SKETCH: The idea behind the proof is illustrated in Figure 1.5. If M uses alphabet Γ then \tilde{M} will use the alphabet Γ^2 (i.e., each symbol in \tilde{M} 's alphabet corresponds to a pair of symbols in M 's alphabet). We encode a tape of M that is infinite in both direction using a standard (infinite in

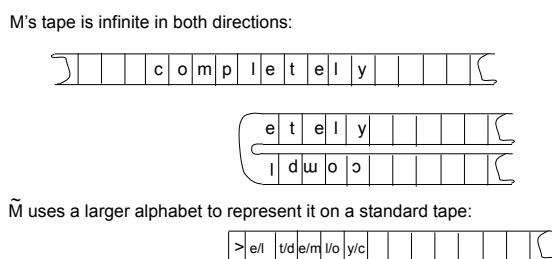


Figure 1.5: To simulate a machine M with alphabet Γ that has tapes infinite in both directions, we use a machine \tilde{M} with alphabet Γ^2 whose tapes encode the “folded” version of M 's tapes.

one direction) tape by “folding” it in an arbitrary location, with each location of \tilde{M} 's tape encoding two locations of M 's tape. At first, \tilde{M} will ignore the second symbol in the cell it reads and act according to M 's transition function. However, if this transition function instructs \tilde{M} to go “over the edge” of its tape then instead it will start ignoring the first symbol in each cell and use only the second symbol. When it is in this mode, it will translate left movements into right movements and vice versa. If it needs to go “over the edge” again then it will go back to reading the first symbol of each cell, and translating movements normally. ■

Other changes that do not have a very significant effect include having two or three dimensional tapes, allowing the machine *random access* to its tape, and making the output tape *write only* (see Exercises 1.11 and 1.13; also the texts [SIP96, HMU01] contain more examples). In particular none of these modifications will change the class **P** of polynomial-time computable decision problems defined below in Section 1.5.

1.3 Machines as strings and the universal Turing machine

It is almost obvious that a Turing machine can be represented as a string: just write the description of the TM on paper, and encode this description as a sequence of zeros and ones. This string can be given as input to another TM—a profound observation that blurs the distinction between computational *hardware*, *software*, and data. It also motivated the invention of the *general purpose* electronic computer, which is a general purpose hardware that can be adapted to any arbitrary task by loading it with an appropriate program (software).

Because we will use this notion of representing TM's as strings quite extensively, it may be worthwhile to spell out our representation a bit more concretely. Since the behavior of a Turing machine is determined by its transition function, we will use the list of all inputs and outputs of this function (which can be easily encoded as a string in $\{0, 1\}^*$) as the encoding of the Turing machine.⁵ We will also find it convenient to assume that our representation scheme satisfies the following properties:

1. Every string in $\{0, 1\}^*$ represents *some* Turing machine.

This is easy to ensure by mapping strings that are not valid encodings into some canonical trivial TM, such as the TM that immediately halts and outputs zero on any input.

2. Every TM is represented by infinitely many strings.

This can be ensured by specifying that the representation can end with an arbitrary number of 1's, that are ignored. This has a somewhat similar effect to the *comments* mechanism of

⁵Note that the size of the alphabet, the number of tapes, and the size of the state space can be deduced from the transition function's table. We can also reorder the table to ensure that the special states $q_{\text{start}}, q_{\text{halt}}$ are the first 2 states of the TM. Similarly, we may assume that the symbols $\triangleright, \square, 0, 1$ are the first 4 symbols of the machine's alphabet.

many programming languages (e.g., the `/*...*/` construct in C, C++ and Java) that allows to add superfluous symbols to any program.

We denote by $\lfloor M \rfloor$ the TM M 's representation as a binary string. If α is a string then M_α denotes the TM that α represents. As is our convention, we will also often use M to denote both the TM and its representation as a string. Exercise 1.15 asks you to fully specify a representation scheme for Turing machines with the above properties.

1.3.1 The Universal Turing Machine

It was Turing who first observed that general purpose computers are possible, by showing a *universal* Turing machine that can *simulate* the execution of every other TM M given M 's description as input. Of course, since we are so used to having a universal computer on our desktops or even in our pockets, today we take this notion for granted. But it is good to remember why it was once counterintuitive. The parameters of the universal TM are fixed —alphabet size, number of states, and number of tapes. The corresponding parameters for the machine being simulated could be much larger. The reason this is not a hurdle is, of course, the ability to use *encodings*. Even if the universal TM has a very simple alphabet, this suffices to represent the other machine's state and transition table on its tapes, and then follow along in the computation step by step.

Now we state a computationally efficient version of Turing's construction due to Hennie and Stearns [HS66]. To give the essential idea we first prove a slightly relaxed variant where the term $T \log T$ below is replaced with T^2 . But since the efficient version is needed a few times in the book, a full proof is also given at the end of the chapter (see Section 1.A).

THEOREM 1.9 (EFFICIENT UNIVERSAL TURING MACHINE)

There exists a TM \mathcal{U} such that for every $x, \alpha \in \{0, 1\}^$, $\mathcal{U}(x, \alpha) = M_\alpha(x)$, where M_α denotes the TM represented by α .*

Furthermore, if M_α halts on input x within T steps then $\mathcal{U}(x, \alpha)$ halts within $CT \log T$ steps, where C is a number independent of $|x|$ and depending only on M_α 's alphabet size, number of tapes, and number of states.

A common exercise in programming courses is to write an *interpreter* for a particular programming language using the same language. (An interpreter takes a program P as input and outputs the result of executing the program P .) Theorem 1.9 can be considered a variant of this exercise.

PROOF OF RELAXED VERSION OF THEOREM 1.9: Our universal TM \mathcal{U} is given an input x, α , where α represents some TM M , and needs to output $M(x)$. A crucial observation is that we may assume that M (**1**) has a single work tape (in addition to the input and output tape) and (**2**) uses the alphabet $\{\triangleright, \square, 0, 1\}$. The reason is that \mathcal{U} can transform a representation of every TM M into a representation of an equivalent TM \tilde{M} that satisfies these properties as shown in the proofs of Claims 1.5 and 1.6. Note that these transformations may introduce a quadratic slowdown (i.e., transform M from running in T time to running in $C'T^2$ time where C' depends on M 's alphabet size and number of tapes).

The TM \mathcal{U} uses the alphabet $\{\triangleright, \square, 0, 1\}$ and three work tapes in addition to its input and output tape (see Figure 1.6). \mathcal{U} uses its input tape, output tape, and one of the work tapes in the same way M uses its three tapes. In addition, \mathcal{U} will use its first extra work tape to store the table of values of M 's transition function (after applying the transformations of Claims 1.5 and 1.6 as noted above), and its other extra work tape to store the current state of M . To simulate one computational step of M , \mathcal{U} scans the table of M 's transition function and the current state to find out the new state, symbols to be written and head movements, which it then executes. We see that

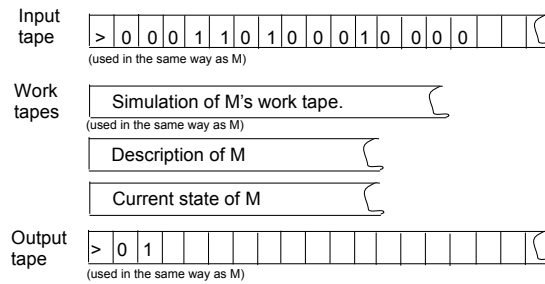


Figure 1.6: The universal TM \mathcal{U} has in addition to the input and output tape, three work tapes. One work tape will have the same contents as the simulated machine M , another tape includes the description M (converted to an equivalent one-work-tape form), and another tape contains the current state of M .

each computational step of M is simulated using C steps of \mathcal{U} , where C is some number depending on the size of the transition function's table.

This high level description can be turned into an exact specification of the TM \mathcal{U} , though we leave this to the reader. If you are not sure how this can be done, think first of how you would program these steps in your favorite programming language and then try to transform this into a description of a Turing machine. ■

It is sometimes useful to consider a variant of the universal TM \mathcal{U} that gets a number T as an extra input (in addition to x and α), and outputs $M_\alpha(x)$ if and only if M_α halts on x within T steps (otherwise outputting some special failure symbol). By adding a counter to \mathcal{U} , the proof of Theorem 1.9 can be easily modified to give such a universal TM.

1.4 Uncomputable functions

It may seem “obvious” that every function can be computed, given sufficient time. However, this turns out to be false: there exist functions that cannot be computed within any finite number of steps!

THEOREM 1.10

There exists a function $\text{UC} : \{0, 1\}^ \rightarrow \{0, 1\}$ that is not computable by any TM.*

PROOF: The function UC is defined as follows: for every $\alpha \in \{0, 1\}^*$, if $M_\alpha(\alpha) = 1$ then $\text{UC}(\alpha) = 0$; otherwise (if $M_\alpha(\alpha)$ outputs a different value or enters an infinite loop), $\text{UC}(\alpha) = 1$.

Suppose for the sake of contradiction that UC is computable and hence there exists a TM M such that $M(\alpha) = \text{UC}(\alpha)$ for every $\alpha \in \{0, 1\}^*$. Then, in particular, $M(\ulcorner M \urcorner) = \text{UC}(\ulcorner M \urcorner)$. But this is impossible: by the definition of UC ,

$$\text{UC}(\ulcorner M \urcorner) = 1 \Leftrightarrow M(\ulcorner M \urcorner) \neq 1.$$

This proof technique is called “diagonalization”: see Figure 1.7. ■

1.4.1 The Halting Problem

One might ask why should we care whether or not the function UC described above is computable—who would want to compute such a contrived function anyway? We now show a more natural uncomputable function. The function **HALT** takes as input a pair $\langle \alpha, x \rangle$ and outputs 1 if and only if the TM M_α represented by α halts on input x within a finite number of steps. This is definitely a function we want to compute: given a computer program and an input we'd certainly like to know if the program is going to enter an infinite loop on this input. Unfortunately, this is not possible, even if we were willing to wait an arbitrarily long time:

	0	1	00	01	10	11	...	α
0	0	1	*	0	1	0		$M_0(\alpha)$	
1	1	1	0	1	*	1		...	
00	*	0	0	1	0	*			
01	1	*	0	0	1	*			
...									
α	$M_\alpha(0)$...						$M_\alpha(\alpha)$	$1 - M_\alpha(\alpha)$
...									

Figure 1.7: Suppose we order all strings in lexicographic order, and write in a table the value of $M_\alpha(x)$ for all strings α, x , where M_α denotes the TM represented by the string α and we use $*$ to denote the case that $M_\alpha(x)$ is not a value in $\{0,1\}$ or that M_α does not halt on input x . Then, UC is defined by “negating” the diagonal of this table, and by its definition it cannot be computed by any TM.

THEOREM 1.11

HALT is not computable by any TM.

PROOF: Suppose, for the sake of contradiction, that there was a TM M_{HALT} computing HALT. We will use M_{HALT} to show a TM M_{UC} computing UC, contradicting Theorem 1.10.

The TM M_{UC} is simple: on input α , M_{UC} runs $M_{\text{HALT}}(\alpha, \alpha)$. If the result is 0 (meaning that M_α does not halt on α) then M_{UC} outputs 1. Otherwise, M_{UC} uses the universal TM \mathcal{U} to compute $b = M_\alpha(\alpha)$. If $b = 1$ M_{UC} outputs 0; otherwise it outputs 1.

Under the assumption that $M_{\text{HALT}}(\alpha, \alpha)$ outputs $\text{HALT}(\alpha, \alpha)$ within a finite number of steps, the TM $M_{\text{UC}}(\alpha)$ will output $\text{UC}(\alpha)$. ■

The proof technique employed to show Theorem 1.11— namely showing that HALT is uncomputable by showing an algorithm for UC using a hypothetical algorithm for HALT— is called a *reduction*. We will see many reductions in this book, often used (as is the case here) to show that a problem B is at least as hard as a problem A by showing an algorithm that could solve A given a procedure that solves B .

There are many other examples for interesting uncomputable (also known as *undecidable*) functions, see Exercise 1.16. There are even uncomputable functions whose formulation has seemingly nothing to do with Turing machines or algorithms. For example, the following problem cannot be solved in finite time by any TM: given a set of polynomial equations with integer coefficients, find out whether these equations have an integer solution (i.e., whether there is an assignment of integers to the variables that satisfies the equations). This is known as the problem of solving Diophantine equations, and in 1900 Hilbert mentioned finding an algorithm for solving it (which he presumed to exist) as one of the top 23 open problems in mathematics.

For more on computability theory, see the chapter notes and the book’s website.

1.5 The class **P**

A *complexity class* is a set of functions that can be computed within given resource bounds. We will now introduce our first complexity classes. For reasons of technical convenience, throughout most of this book we will pay special attention to Boolean functions (that have one bit output), also known as *decision problems* or *languages*.

DEFINITION 1.12 (THE CLASS **DTIME**.) □

Let $T : \mathbb{N} \rightarrow \mathbb{N}$ be some function. We let $\mathbf{DTIME}(T(n))$ be the set of all Boolean (one bit output) functions that are computable in $c \cdot T(n)$ -time for some constant $c > 0$.

The “D” in the notation **DTIME** refers to “deterministic. The Turing machine introduced in this chapter is more precisely called the *deterministic* Turing machine since for any given input x , it can compute in exactly one way. Later we will see other types of Turing machines, including nondeterministic and probabilistic TMs.

The following class will serve as our rough approximation for the class of decision problems that are efficiently solvable.

DEFINITION 1.13 (THE CLASS **P**)

$$\mathbf{P} = \cup_{c \geq 1} \mathbf{DTIME}(n^c)$$

Thus, we can phrase the question from the introduction as to whether the dinner party problem has an efficient algorithm as follows: “Is **INDSET** in **P**?”, where **INDSET** is the language defined in Example 1.1.

EXAMPLE 1.14

We give some examples to emphasize a couple of points about the definition of the class **P**. First, the class contains only decision problems. Thus we cannot say, for example, that “Integer multiplication is in **P**.” Instead, we may say that its decision version is in **P**, namely, the following language:

$$\left\{ \langle x, i \rangle : \text{The } i^{\text{th}} \text{ bit of } xy \text{ is equal to } 1 \right\} .$$

Second, the running time is a function of the number of *bits* in the input. Consider the problem of solving a system of linear equations over the rational numbers. In other words, given a pair $\langle A, \mathbf{b} \rangle$ where A is an $m \times n$ rational matrix and \mathbf{b} is an m dimensional rational vector, find out if there exists an n -dimensional vector \mathbf{x} such that $A\mathbf{x} = \mathbf{b}$. The standard Gaussian Elimination algorithm solves this problem in $O(n^3)$ *arithmetic operations*. But on a Turing machine, each arithmetic operation has to be done in the gradeschool fashion, bit by laborious bit. Thus to prove that this decision problem is in **P** we have to verify that Gaussian elimination (or some other algorithm for the problem) runs on a Turing machine in time that is polynomial in the number of bits required to represent a_1, a_2, \dots, a_n . That is, in the case of Gaussian elimination, we need to verify that all the intermediate numbers involved in the computation can be represented by polynomially many bits. Fortunately, this does turn out to be the case (for a related result, see Exercise 2.3).

1.5.1 On the philosophical importance of **P**

The class **P** is felt to capture the notion of decision problems with “feasible” decision procedures. Of course, one may argue whether $\mathbf{DTIME}(n^{100})$ really represents “feasible” computation in the

real world since n^{100} is prohibitively huge even for moderate values of n . However, in practice, whenever we show that a problem is in \mathbf{P} , we usually find an n^3 or n^5 time algorithm (with reasonable constants), and not an n^{100} time algorithm. (It has also happened a few times that the first polynomial-time algorithm for a problem had high complexity, say n^{20} , but soon somebody simplified it to say an n^5 time algorithm.)

Note that the class \mathbf{P} is useful only in a certain context. Turing machines are a crude model if one is designing algorithms that must run in a fraction of a second on the latest PC (in which case one must carefully account for fine details about the hardware). However, if the question is whether any subexponential algorithms exist for, say, the language INDSET of Example 0.1, then even an n^{20} time algorithm would be a fantastic breakthrough.

\mathbf{P} is also a natural class from the viewpoint of a programmer. Suppose undergraduate programmers are asked to invent the definition of an “efficient” computation. Presumably, they would agree that a computation that runs in linear or quadratic time is “efficient.” Next, since programmers often write programs that call other programs (or subroutines), they might find it natural to consider a program “efficient” if it performs only “efficient” computations and calls subroutines that are “efficient”. The resulting notion of “efficient computations” obtained turns out to be exactly the class \mathbf{P} [Cob64].

1.5.2 Criticisms of \mathbf{P} and some efforts to address them

Now we address some possible criticisms of the definition of \mathbf{P} , and some related complexity classes that address these.

Worst-case exact computation is too strict. The definition of \mathbf{P} only considers algorithms that compute the function on *every* possible input. However, not all possible inputs arise in practice (although it’s not always easy to characterize the inputs that do). Chapter 17 gives a theoretical treatment of *average-case complexity* and defines the analogue of \mathbf{P} in that context. Sometimes, users are willing to settle for *approximate* solutions. Chapter 21 contains a rigorous treatment of the complexity of approximation.

Other physically realizable models. If we were to make contact with an advanced alien civilization, with a different computational model than the Turing machine, would their class \mathbf{P} be any different from the class defined here?

Most scientists believe the *Church-Turing (CT) thesis*, which states that every physically realizable computation device— whether it’s silicon-based, DNA-based, neuron-based or using some alien technology— can be simulated by a Turing machine. Thus they believe that the set of *computable* problems would be the same for aliens as it is for us. (The CT thesis is not a theorem, merely a belief about the nature of the world.)

However, when it comes to *efficiently* computable problems, the situation is less clear. The **strong form of the CT thesis** says that every physically realizable computation model can be simulated by a TM *with polynomial overhead* (in other words, t steps on the model can be simulated in t^c steps on the TM, where c is a constant that depends upon the model). If true, it implies that the class \mathbf{P} defined by the aliens will be the same as ours. However, several objections have been made to this strong form.

(a) *Issue of precision:* TMs compute with discrete symbols, whereas physical quantities may be real numbers in \mathbb{R} . Thus TM computations may only be able to approximately simulate the real world. Though this issue is not perfectly settled, it seems so far that TMs do not suffer from an inherent handicap. After all, real-life devices suffer from noise, and physical quantities can only be measured up to finite precision. Thus a TM could simulate the real-life device using finite precision. (Note also that we often only care about the *most significant bit* of the result, namely, a 0/1 answer.)

Even so, in Chapter 15 we also consider a modification of the TM model that allows computations in \mathbb{R} as a basic operation. The resulting complexity classes have fascinating connections with the usual complexity classes.

(b) *Use of randomness:* The TM as defined is *deterministic*. If randomness exists in the world, one can conceive of computational models that use a source of random bits (i.e., "coin tosses"). Chapter 7 considers Turing Machines that are allowed to also toss coins, and studies the complexity class \mathbf{BPP} , that is the analogue of \mathbf{P} for those machines. (However, we will see in Chapters 18 and 19 the intriguing possibility that randomized computation may be no more powerful than deterministic computation.)

(c) *Use of quantum mechanics:* A more clever computational model might use some of the counterintuitive features of quantum mechanics. In Chapter 10 we define the complexity class \mathbf{BQP} , that generalizes \mathbf{P} in such a way. We will see problems in \mathbf{BQP} that are currently not known to be in \mathbf{P} (though there is no known proof that $\mathbf{BQP} \neq \mathbf{P}$). However, it is not yet clear whether the quantum model is truly physically realizable. Also quantum computers currently seem only able to efficiently solve only very few "well-structured" problems that are (presumed to be) not in \mathbf{P} . Hence some insights gained from studying \mathbf{P} may still apply to quantum computers.

(d) *Use of other exotic physics, such as string theory.* Though an intriguing possibility, it hasn't yet had the same scrutiny as use of quantum mechanics.

Decision problems are too limited. Some computational problems are not easily expressed as decision problems. Indeed, we will introduce several classes in the book to capture tasks such as computing non-Boolean functions, solving search problems, approximating optimization problems, interaction, and more. Yet the framework of decision problems turn out to be surprisingly expressive, and we will often use it in this book.

1.5.3 Edmonds' quote

We conclude this section with a quote from Edmonds [Edm65], that in the paper showing a polynomial-time algorithm for the maximum matching problem, explained the meaning of such a result as follows:

For practical purposes computational details are vital. However, my purpose is only to show as attractively as I can that there is an efficient algorithm. According to the dictionary, "efficient" means "adequate in operation or performance." This is roughly the meaning I want in the sense that it is conceivable for maximum matching to have no efficient algorithm.

...There is an obvious finite algorithm, but that algorithm increases in difficulty exponentially with the size of the graph. It is by no means obvious whether or not there exists an algorithm whose difficulty increases only algebraically with the size of the graph.

...When the measure of problem-size is reasonable and when the sizes assume values arbitrarily large, an asymptotic estimate of ... the order of difficulty of an algorithm is theoretically important. It cannot be rigged by making the algorithm artificially difficult for smaller sizes.

...One can find many classes of problems, besides maximum matching and its generalizations, which have algorithms of exponential order but seemingly none better ... For practical purposes the difference between algebraic and exponential order is often more crucial than the difference between finite and non-finite.

...It would be unfortunate for any rigid criterion to inhibit the practical development of algorithms which are either not known or known not to conform nicely to the criterion.

Many of the best algorithmic idea known today would suffer by such theoretical pedantry. ... However, if only to motivate the search for good, practical algorithms, it is important to realize that it is mathematically sensible even to question their existence. For one thing the task can then be described in terms of concrete conjectures.

WHAT HAVE WE LEARNED?

- There are many equivalent ways to mathematically model computational processes; we use the standard Turing machine formalization.
- Turing machines can be represented as strings. There is a *universal* TM that can emulate (with small overhead) any TM given its representation.
- There exist functions, such as the Halting problem, that cannot be computed by any TM regardless of its running time.
- The class \mathbf{P} consists of all decision problems that are solvable by Turing machines in polynomial time. We say that problems in \mathbf{P} are efficiently solvable.
- All low-level choices (number of tapes, alphabet size, etc..) in the definition of Turing machines are immaterial, as they will not change the definition of \mathbf{P} .

Chapter notes and history

Although certain algorithms have been studied for thousands of years, and some forms of computing devices were designed before the 20th century (most notably Charles Babbage’s difference and analytical engines in the mid 1800’s), it seems fair to say that the foundations of modern computer science were only laid in the 1930’s.

In 1931, Kurt Gödel shocked the mathematical world by showing that certain true statements about the natural numbers are *inherently unprovable*, thereby shattering an ambitious agenda set in 1900 by David Hilbert to base all of mathematics on solid axiomatic foundations. In 1936, Alonzo Church defined a model of computation called λ -calculus (which years later inspired the programming language LISP) and showed the existence of functions *inherently uncomputable* in this model [Chu36]. A few months later, Alan Turing independently introduced his Turing machines and showed functions inherently uncomputable by such machines [Tur36]. Turing also introduced the idea of the *universal* Turing machine that can be loaded with arbitrary programs. The two models turned out to be equivalent, but in the words of Church himself, Turing machines have “the advantage of making the identification with effectiveness in the ordinary (not explicitly defined) sense evident immediately”. The anthology [Dav65] contains many of the seminal papers in the theory of computability. Part II of Sipser’s book [SIP96] is a good gentle introduction to this theory, while the books [rogers, HMU01, Koz97] go into a bit more depth. This book’s web site also contains some additional coverage of this theory.

During World War II Turing designed mechanical code-breaking devices and played a key role in the effort to crack the German “Enigma” cipher, an achievement that had a decisive effect on the war’s progress (see the biographies [Hod83, Lea05]).⁶ After World War II, efforts to build electronic universal computers were undertaken in both sides of the Atlantic. A key figure in these efforts was John von-Neumann, an extremely prolific scientist that was involved in everything from the Manhattan project to founding game theory in economics. To this day essentially all digital

⁶Unfortunately, Turing’s wartime achievements were kept confidential during his lifetime, and so did not keep him from being forced by British courts to take hormones to “cure” his homosexuality, resulting in his suicide in 1954.

computers follow the “von-Neumann architecture” he pioneered while working on the design of the EDVAC, one of the earliest digital computers [vN45].

As computers became more prevalent, the issue of efficiency in computation began to take center stage. Cobham [Cob64] defined the class **P** and suggested it may be a good formalization for efficient computation. A similar suggestion was made by Edmonds ([Edm65], see quote above) in the context of presenting a highly non-trivial polynomial-time algorithm for finding a maximum matching in general graphs. Hartmanis and Stearns [HS65] defined the class **DTIME**($T(n)$) for every function T , and proved the slightly relaxed version of Theorem 1.9 we showed above (the version we stated and prove below was given by Hennie and Stearns [HS66]). They also coined the name “computational complexity” and proved an interesting “speed-up theorem”: if a function f is computable by a TM M in time $T(n)$ then for every constant $c \geq 1$, f is computable by a TM \tilde{M} (possibly with larger state size and alphabet size than M) in time $T(n)/c$. This speed-up theorem is another justification for ignoring constant factors in the definition of **DTIME**($T(n)$). Blum [Blu67] suggested an axiomatic formalization of complexity theory, that does not explicitly mention Turing machines.

We have omitted a discussion of some of the “bizarre conditions” that may occur when considering time bounds that are not time-constructible, especially “huge” time bounds (i.e., function $T(n)$ that are much larger than exponential in n). For example, there is a non-time constructible function $T : \mathbb{N} \rightarrow \mathbb{N}$ such that every function computable in time $T(n)$ can also be computed in the much shorter time $\log T(n)$. However, we will not encounter non time-constructible time bounds in this book.

The result that **PAL** requires $\Omega(n^2)$ steps to compute on TM’s using a single read/write tape is from [??]. A proof can be found in [??].

We have stated that algorithms that take less than n steps are not very interesting as they do not even have time to read their input. This is true for the Turing machine model. However, if one allows *random access* to the input combined with *randomization* then many interesting computational tasks can actually be achieved in *sublinear* time. See [??] for a survey of this line of research.

Exercises

- 1.1. For each of the following pairs of functions f, g determine whether $f = o(g)$, $g = o(f)$ or $f = \Theta(g)$. If $f = o(g)$ then find the first number n such that $f(n) < g(n)$:

(a) $f(n) = n^2$, $g(n) = 2n^2 + 100\sqrt{n}$.

(b) $f(n) = n^{100}$, $g(n) = 2^{n/100}$.

(c) $f(n) = n^{100}$, $g(n) = 2^{n^{1/100}}$.

(d) $f(n) = \sqrt{n}$, $g(n) = 2^{\sqrt{\log n}}$.

(e) $f(n) = n^{100}$, $g(n) = 2^{(\log n)^2}$.

(f) $f(n) = 1000n$, $g(n) = n \log n$.

- 1.2. For each of the following recursively defined functions f , find a closed (non-recursive) expression for a function g such that $f(n) = \Theta(g(n))$, and prove that this is the case.

(Note: below we only supply the recursive rule, you can assume that $f(1) = f(2) = \dots = f(10) = 1$ and the recursive rule is applied for $n > 10$; in any case regardless of how the base case is defined it won’t make any difference to the answer - can you see why?)

(a) $f(n) = f(n - 1) + 10$.

(b) $f(n) = f(n - 1) + n$.

- (c) $f(n) = 2f(n - 1)$.
- (d) $f(n) = f(n/2) + 10$.
- (e) $f(n) = f(n/2) + n$.
- (f) $f(n) = 2f(n/2) + n$.
- (g) $f(n) = 3f(n/2)$.
- (h) $f(n) = 2f(n/2) + O(n^2)$.

Hint: Answers are: (a) n (b) n^2 (c) 2^n (d) $\log n$ (e) n (f) $n \log n$ (g) $n^{\log 3}$ (h) n^2 .

- 1.3. The MIT museum contains a kinetic sculpture by Arthur Ganson called “Machine with concrete” (see Figure 1.8). It consists of 13 gears connected to one another in a series such that each gear moves 50 times slower than the previous one. The fastest gear is constantly rotated by an engine at a rate of 212 rotations per minute. The slowest gear is fixed to a block of concrete and so apparently cannot move at all. How come this machine does not break apart?




Figure 1.8: Machine with concrete by Arthur Ganson.

- 1.4. Let f be the *addition* function that maps the representation of a pair of numbers x, y to the representation of the number $x + y$. Let g be the *multiplication* function that maps $\langle x, y \rangle$ to $\lfloor x \cdot y \rfloor$. Prove that both f and g are computable by writing down a full description (including the states, alphabet and transition function) of the corresponding Turing machines.

Hint: Follow the gradeschool algorithms.

- 1.5. Complete the proof of Claim 1.5 by writing down explicitly the description of the machine \tilde{M} .
- 1.6. Complete the proof of Claim 1.6.
- 1.7. Complete the proof of Claim 1.8.
- 1.8. Define a TM M to be *oblivious* if its head movements do not depend on the input but only on the input length. That is, M is oblivious if for every input $x \in \{0, 1\}^*$ and $i \in \mathbb{N}$, the location of each of M 's heads at the i^{th} step of execution on input x is only a function of $|x|$ and i . Show that for every time-constructible $T : \mathbb{N} \rightarrow \mathbb{N}$, if $L \in \mathbf{DTIME}(T(n))$ then there is an oblivious TM that decides L in time $O(T(n)^2)$. Furthermore, show that there is such a TM that uses only *two tapes*: one input tape and one work/output tape.

Hint: Use the proof of Claim 1.6.

- 1.9. Show that for every time-constructible $T : \mathbb{N} \rightarrow \mathbb{N}$, if $L \in \mathbf{DTIME}(T(n))$ then there is an oblivious TM that decides L in time $O(T(n) \log T(n))$.

Hint: show that the universal TM \mathcal{U} obtained by the proof of Theorem 1.9 can be tweaked to be oblivious.

- 1.10. Define a *single-tape* Turing machine to be a TM that has only one read/write tape, that is used as input, work and output tape. Show that for every (time-constructible) $T : \mathbb{N} \rightarrow \mathbb{N}$ and $f \in \mathbf{DTIME}(T(n))$, f can be computed in $O(T(n)^2)$ steps by a single-tape TM.
- 1.11. Define a *two dimensional* Turing machine to be a TM where each of its tapes is an infinite grid (and the machine can move not only Left and Right but also Up and Down). Show that for every (time-constructible) $T : \mathbb{N} \rightarrow \mathbb{N}$ and every Boolean function f , if g can be computed in time $T(n)$ using a two-dimensional TM then $f \in \mathbf{DTIME}(T(n)^2)$.
- 1.12. Let LOOKUP denote the following function: on input a pair $\langle x, i \rangle$ (where x is a binary string and i is a natural number), LOOKUP outputs the i^{th} bit of x or 0 if $|x| < i$. Prove that LOOKUP $\in \mathbf{P}$.
- 1.13. Define a *RAM Turing machine* to be a Turing machine that has *random access memory*. We formalize this as follows: the machine has additional two symbol on its alphabet we denote by R and W and an additional state we denote by q_{access} . We also assume that the machine has an infinite array A that is initialized to all blanks. Whenever the machine enters q_{access} , if its address tape contains $\ulcorner i \urcorner R$ (where $\ulcorner i \urcorner$ denotes the binary representation of i) then the value $A[i]$ is written in the cell next to the R symbol. If its tape contains $\ulcorner i \urcorner W \sigma$ (where σ is some symbol in the machine's alphabet) then $A[i]$ is set to the value σ .
- Show that if a Boolean function f is computable within time $T(n)$ (for some time-constructible T) by a RAM TM, then it is in $\mathbf{DTIME}(T(n)^2)$.
- 1.14. Consider the following simple programming language. It has a single infinite array A of elements in $\{0, 1, \square\}$ (initialized to \square) and a single integer variable i . A program in this language contains a sequence of lines of the following form:

label : If $A[i]$ equals σ then *cmds*

Where $\sigma \in \{0, 1, \square\}$ and *cmds* is a list of one or more of the following commands: **(1) Set $A[i]$ to τ** where $\tau \in \{0, 1, \square\}$, **(2) Goto *label***, **(3) Increment i by one**, **(4) Decrement i by one**, and **(5) Output b and halt.** where $b \in \{0, 1\}$. A program is executed on an input $x \in \{0, 1\}^n$ by placing the i^{th} bit of x in $A[i]$ and then running the program following the obvious semantics.

Prove that for every functions $f : \{0, 1\}^* \rightarrow \{0, 1\}$ and (time constructible) $T : \mathbb{N} \rightarrow \mathbb{N}$, if f is computable in time $T(n)$ by a program in this language, then $f \in \mathbf{DTIME}(T(n))$.

- 1.15. Give a full specification of a representation scheme of Turing machines as binary string strings. That is, show a procedure that transforms any TM M (e.g., the TM computing the function PAL described in Example 1.1) into a binary string $\ulcorner M \urcorner$. It should be possible to recover M from $\ulcorner M \urcorner$, or at least recover a functionally equivalent TM (i.e., a TM \tilde{M} computing the same function as M with the same running time).
- 1.16. A *partial* function from $\{0, 1\}^*$ to $\{0, 1\}^*$ is a function that is not necessarily defined on all its inputs. We say that a TM M computes a partial function f if for every x on which f is defined, $M(x) = f(x)$ and for every x on which f is not defined M gets into an infinite loop when executed on input x . If \mathcal{S} is a set of partial functions, we define $f_{\mathcal{S}}$ to be the Boolean function that on input α outputs 1 iff M_{α} computes a partial function in \mathcal{S} . *Rice's Theorem* says that for every non-trivial \mathcal{S} (a set that is not the empty set nor the set of all partial functions), the $f_{\mathcal{S}}$ is not computable.

- (a) Show that Rice's Theorem yields an alternative proof for Theorem 1.11 by showing that the function HALT is not computable.
- (b) Prove Rice's Theorem.

Hint: By possibly changing from \mathcal{S} to its complement, we may assume that the empty function \emptyset (that is not defined on any input) is in \mathcal{S} there is some function f that is defined on some input x that is not in \mathcal{S} . Use this to show that an algorithm to compute $f_{\mathcal{S}}$ can compute the function HALT_x which outputs 1 on input α iff M_{α} halts on input x . Then reduce computing HALT to computing HALT_x thereby deriving Rice's Theorem from Theorem 1.11.

1.17. Prove that the following languages/decision problems on graphs are in \mathbf{P} : (You may pick either the adjacency matrix or adjacency list representation for graphs; it will not make a difference. Can you see why?)

- (a) CONNECTED — the set of all connected graphs. That is, $G \in \text{CONNECTED}$ if every pair of vertices u, v in G are connected by a path.
- (b) TRIANGLEFREE — the set of all graphs that do not contain a triangle (i.e., a triplet u, v, w of connected distinct vertices).
- (c) BIPARTITE — the set of all bipartite graphs. That is, $G \in \text{BIPARTITE}$ if the vertices of G can be partitioned to two sets A, B such that all edges in G are from a vertex in A to a vertex in B (there is no edge between two members of A or two members of B).
- (d) TREE — the set of all trees. A graph is a *tree* if it is connected and contains no cycles. Equivalently, a graph G is a tree if every two distinct vertices u, v in G are connected by exactly one simple path (a path is simple if it has no repeated vertices).

1.18. Recall that normally we assume that numbers are represented as string using the *binary* basis. That is, a number n is represented by the sequence $x_0, x_1, \dots, x_{\log n}$ such that $n = \sum_{i=0}^n x_i 2^i$. However, we could have used other encoding schemes. If $n \in \mathbb{N}$ and $b \geq 2$, then the *representation of n in base b* , denoted by $\llcorner n \lrcorner_b$ is obtained as follows: first represent n as a sequence of digits in $\{0, \dots, b-1\}$, and then replace each digit $d \in \{0, \dots, b-1\}$ by its binary representation. The *unary* representation of n , denoted by $\llcorner n \lrcorner_{\text{unary}}$ is the string 1^n (i.e., a sequence of n ones).

- (a) Show that choosing a different base of representation will make no difference to the class \mathbf{P} . That is, show that for every subset S of the natural numbers, if we define $L_S^b = \{ \llcorner n \lrcorner_b : n \in S \}$ then for every $b \geq 2$, $L_S^b \in \mathbf{P}$ iff $L_S^2 \in \mathbf{P}$.
- (b) Show that choosing the unary representation may make a difference by showing that the following language is in \mathbf{P} :

$$\text{UNARYFACTORING} = \{ \langle \llcorner n \lrcorner_{\text{unary}}, \llcorner \ell \lrcorner_{\text{unary}}, \llcorner k \lrcorner_{\text{unary}} \rangle : \text{there is a prime } j \in (\ell, k) \text{ dividing } n \}$$

It is not known to be in \mathbf{P} if we choose the binary representation (see Chapters 9 and 10). In Chapter 3 we will see that there is a problem that is *proven* to be in \mathbf{P} when choosing the unary representation but not in \mathbf{P} when using the binary representation.

1.A Proof of Theorem 1.9: Universal Simulation in $O(T \log T)$ -time

We now show how to prove Theorem 1.9 as stated. That is, we show a universal TM \mathcal{U} such that given an input x and a description of a TM M that halts on x within T steps, \mathcal{U} outputs

$M(x)$ within $O(T \log T)$ time (where the constants hidden in the O notation may depend on the parameters of the TM M being simulated).

The general structure of \mathcal{U} will be as in Section 1.3.1. \mathcal{U} will use its input and output tape in the same way M does, and will also have extra work tapes to store M 's transition table and current state, and to encode the contents of M 's work tapes. The main obstacle we need to overcome is that we cannot use Claim 1.6 to reduce the number of M 's work tapes to one, since Claim 1.6 introduces too much overhead in the simulation. Therefore, we will show a different way to encode all of M 's work tapes in a single tape of \mathcal{U} , which we call the *main* work tape of \mathcal{U} .

Let k be the number of tapes that M uses (apart from its input and output tapes) and Γ its alphabet. Following the proof of Claim 1.5, we may assume that \mathcal{U} uses the alphabet Γ^k (as this can be simulated with a overhead depending only on $k, |\Gamma|$). Thus we can encode in each cell of \mathcal{U} 's main work tape k symbols of Γ , each corresponding to a symbol from one of M 's tapes. This means that we can think of \mathcal{U} 's main work tape not as a single tape but rather as k *parallel tapes*; that is, we can think of \mathcal{U} as having k tapes with the property that in each step all their read/write heads go in unison either one location to the left, one location to the right or they all stay in place. While we can easily encode the contents of M 's k work tapes in \mathcal{U} 's k parallel tapes, we still have to deal with the fact that M 's k read/write heads can each move independently to the left or right, whereas \mathcal{U} 's parallel tapes are forced to move together. Paraphrasing the famous saying, our strategy to handle this is: “*If the head cannot go to the tape locations then the locations will go to the head*”.

That is, since we can not move \mathcal{U} 's read/write head in different directions at once, we simply move the parallel tapes “under” the head. To simulate a single step of M we shift all the non-blank symbols in each of these parallel tapes until the head's position in these parallel tapes corresponds to the heads' positions of M 's k tapes. For example, if $k = 3$ and in some particular step M 's transition function specifies the movements L, R, R then \mathcal{U} will shift all the non-blank entries of its first parallel tape one cell to the right, and shift the non-blank entries of its second and third tapes one cell to the left (see Figure 1.9). \mathcal{U} can easily perform such shifts using an additional “scratch” work tape.




Figure 1.9: Packing k tapes of M into one tape of \mathcal{U} . We consider \mathcal{U} 's single work tape to be composed of k parallel tapes, whose heads move in unison, and hence we shift the contents of these tapes to simulate independent head movement.

The approach above is still not good enough to get $O(T \log T)$ -time simulation. The reason is that there may be as many as T non-blank symbols in each parallel tape, and so each shift operation may cost \mathcal{U} as much as T operations per each step of M , resulting in $\Theta(T^2)$ -time simulation. We will deal with this problem by encoding the information on the tapes in a way that allows us to amortize the work of performing a shift. We will ensure that we do not need to move the entire non-blank symbols of the tape in each shift operation. Specifically, we will encode the information in a way that allows half of the shift operations to be performed using $2c$ steps, for some constant c , a quarter of them using $4c$ steps, and more generally 2^{-i} fraction of the operations will take $2^i c$ steps, leading to simulation in roughly $cT \log T$ time (see below). (This kind of analysis is called *amortized analysis* and is widely used in algorithm design.)

Encoding M 's tapes on \mathcal{U} 's tape. To allow more efficient shifts we encode the information using “buffer zones”: rather than having each of \mathcal{U} 's parallel tapes correspond exactly to a tape of M , we add a special kind of blank symbol \boxtimes to the alphabet of \mathcal{U} 's parallel tapes with the semantics that this symbol is ignored in the simulation. For example, if the non-blank contents of M 's tape are 010 then this can be encoded in the corresponding parallel tape of \mathcal{U} not just by 010 but also by $0\boxtimes 01$ or $0\boxtimes\boxtimes 1\boxtimes 0$ and so on..

For convenience, we think of \mathcal{U} 's parallel tapes as infinite in both the left and right directions (this can be easily simulated with minimal overhead: see Claim 1.8). Thus, we index their locations by $0, \pm 1, \pm 2, \dots$. Normally we keep \mathcal{U} 's head on location 0 of these parallel tapes. We will only move it temporarily to perform a shift when, following our general approach, we simulate a left head movement by shifting the tape to the right and vice versa. At the end of the shift we return the head to location 0.

We split each of \mathcal{U} 's parallel tapes into *zones* that we denote by $R_0, L_0, R_1, L_1, \dots$ (we'll only need to go up to $R_{\log T}, L_{\log T}$). The cell at location 0 is not at any zone. Zone R_0 contains the two cells immediately to the right of location C (i.e., locations $+1$ and $+2$), while Zone R_1 contains the four cells $+3, +4, +5, +6$. Generally, for every $i \geq 1$, Zone R_i contains the $2 \cdot 2^i$ cells that are to the right of Zone R_{i-1} (i.e., locations $[2^{i+1} - 1, 2^{i+2} - 2]$). Similarly, Zone L_0 contains the two cells indexed by -1 and -2 , and generally Zone L_i contains the cells $[-2^{i+2} + 2, -2^{i+1} + 1]$. We shall always maintain the following invariants:

- Each of the zones is either *empty*, *full*, or *half-full* with non- \boxtimes symbols. That is, the number of symbols in zone R_i that are not \boxtimes is either $0, 2^i$, or $2 \cdot 2^i$ and the same holds for L_i . (We treat the ordinary \square symbol the same as any other symbol in Γ and in particular a zone full of \square 's is considered full.)

We assume that initially all the zones are half-full. We can ensure this by filling half of each zone with \boxtimes symbols in the first time we encounter it.

- The total number of non- \boxtimes symbols in $R_i \cup L_i$ is $2 \cdot 2^i$. That is, either R_i is empty and L_i is full, or R_i is full and L_i is empty, or they are both half-full.
- Location 0 always contains a non- \boxtimes symbol.

Performing a shift. The advantage in setting up these zones is that now when performing the shifts, we do not always have to move the entire tape, but we can restrict ourselves to only using some of the zones. We illustrate this by showing how \mathcal{U} performs a left shift on the first of its parallel tapes (see also Figure 1.10):

1. \mathcal{U} finds the smallest i_0 such that R_{i_0} is not empty. Note that this is also the smallest i_0 such that L_{i_0} is not full. We call this number i_0 the *index* of this particular shift.

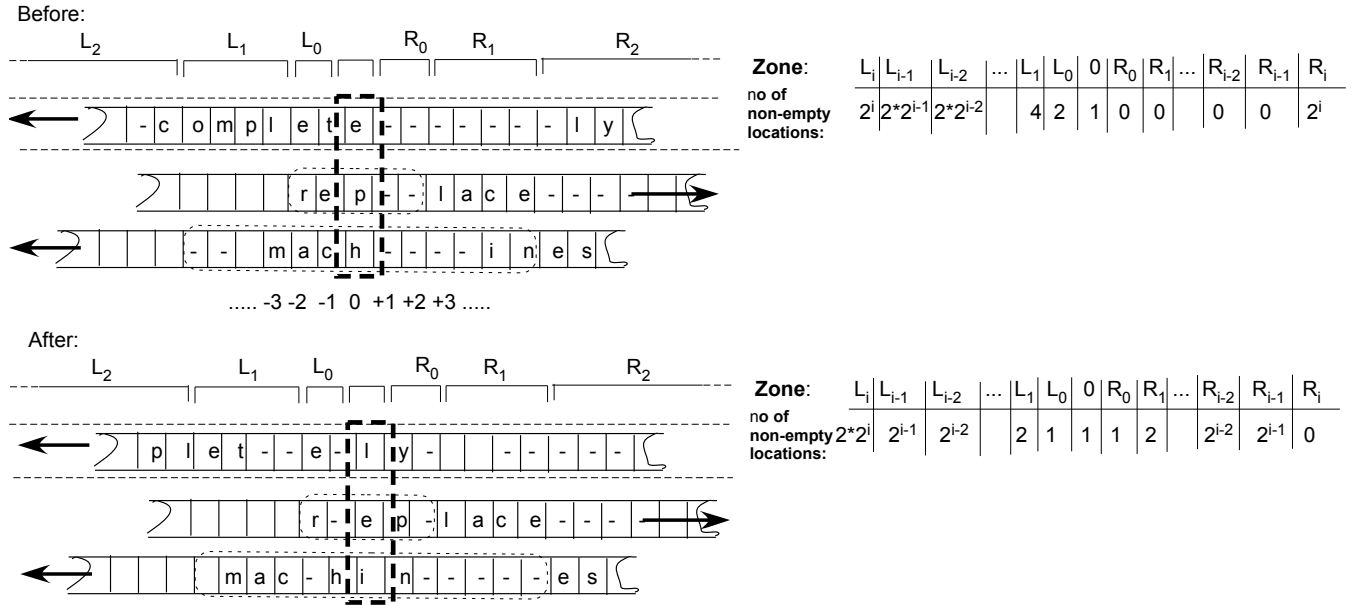


Figure 1.10: Performing a shift of the parallel tapes. The left shift of the first tape involves zones $R_0, L_0, R_1, L_1, R_2, L_2$, the right shift of the second tape involves only R_0, L_0 , while the left shift of the third tape involves zones R_0, L_0, R_1, L_1 . We maintain the invariant that each zone is either empty, half-full or full and that the total number of non-empty cells in $R_i \cup L_i$ is $2 \cdot 2^i$. If before the left shift zones L_0, \dots, L_{i-1} were full and L_i was half-full (and so R_0, \dots, R_{i-1} were full and R_i half-full), then after the shift zones $R_0, L_0, \dots, R_{i-1}, L_{i-1}$ will be half-full, L_i will be full and R_i will be empty. (In this figure - denotes \square .)

- \mathcal{U} puts the leftmost non- \square symbol of R_{i_0} in position 0 and shifts the remaining leftmost $2^{i_0} - 1$ non- \square symbols from R_{i_0} into the zones R_0, \dots, R_{i_0-1} filling up exactly half the symbols of each zone. Note that there is exactly room to perform this since all the zones R_0, \dots, R_{i_0-1} were empty and indeed $2^{i_0} - 1 = \sum_{j=0}^{i_0-1} 2^j$.
- \mathcal{U} performs the symmetric operation to the left of position 0. That is, for $j = i_0 - 1..0$ it iteratively moves the $2 \cdot 2^j$ symbols from L_j to fill half the cells of L_{j+1} . Finally, \mathcal{U} moves the symbol originally in position 0 (modified appropriately according to M 's transition function) to L_0 .
- At the end of the shift, all of the zones $R_0, L_0, \dots, R_{i_0-1}, L_{i_0-1}$ are half-full, R_{i_0} has 2^{i_0} fewer non- \square symbols, and L_i has 2^i additional non- \square symbols. Thus, our invariants are maintained.
- The total cost of performing the shift is proportional to the total size of all the zones involved $R_0, L_0, \dots, R_{i_0}, L_{i_0}$. That is, $O(\sum_{j=0}^{i_0} 2 \cdot 2^j) = O(2^{i_0})$ operations.

After performing a shift with index i the zones $L_0, R_0, \dots, L_{i-1}, R_{i-1}$ are half-full, which means that it will take at least $2^i - 1$ left shifts before the zones L_0, \dots, L_{i-1} become empty or at least $2^i - 1$ right shifts before the zones R_0, \dots, R_{i-1} become empty. In any case, once we perform a shift with index i , the next $2^i - 1$ shifts of that particular parallel tape will all have index less than i . This means that for every one of the parallel tapes, at most a $1/2^i$ fraction of the total number of shifts have index i . Since we perform at most T shifts, and the highest possible index is $\log T$, the total work spent in shifting \mathcal{U} 's k parallel tapes in the course of simulating T steps of M is

$$O(k \cdot \sum_{i=1}^{\log T} \frac{T}{2^{i-1}} 2^i) = O(T \log T).$$

■

