

Computability and Computational Complexity  
Academic year 2018–2019, first semester  
Lecture notes

Mauro Brunato

Version: 2018-11-05

### **Warning**

The main purpose of these very schematic lecture notes is to keep track of what has been said during the lectures. Reading this document is not enough to pass the exam. You should also see the linked citations and footnotes, the additional material and the references provided on the following webpage, which will also contain the up-to-date version of these notes:

<https://comp3.eu/>

To check for new version of this document, please compare the version date on the title page the one reported on the webpage.

# Changelog

## 2018-11-05

- Computational complexity: introductory part (**P**, **NP**, reductions).

## 2018-10-10

- Answers to exercises.

## 2018-10-08

- Rice's theorem, non-computability of the BB function.

## 2018-10-05

- Added some nomenclature and exercises.

## 2018-10-01

- Fourth and fifth lecture: existence of uncomputable functions, the halting problem, co-R.E. functions and sets, the Busy Beaver game.

## 2018-09-24

- Third lecture: alphabet size, Universal Turing machines, emulation of other computational systems.

## 2018-09-19

- Second lecture: definition of Turing machine, examples, equivalence of multiple-tape machines.

## 2018-09-14

- First lecture: basic definitions, Collatz example

# Contents

<b>I</b>	<b>Lecture notes</b>	<b>3</b>
<b>1</b>	<b>Computability</b>	<b>4</b>
1.1	Basic definitions and examples . . . . .	4
1.1.1	A possibly non-recursive set . . . . .	5
1.2	A computational model: the Turing machine . . . . .	6
1.2.1	Examples . . . . .	8
1.2.2	Computational power of the Turing Machine . . . . .	8
1.2.3	Universal Turing machines . . . . .	9
1.2.4	The Church-Turing thesis . . . . .	10
1.3	Uncomputable functions . . . . .	10
1.3.1	Finding an uncomputable function . . . . .	11
1.3.2	Recursive enumerability of halting computations . . . . .	13
1.3.3	Another uncomputable function: the Busy Beaver game . . . . .	14
1.3.4	More definitions . . . . .	16
1.3.5	Rice's Theorem . . . . .	16
<b>2</b>	<b>Complexity classes: P and NP</b>	<b>18</b>
2.1	Definitions . . . . .	18
2.2	Polynomial languages . . . . .	19
2.2.1	Examples . . . . .	19
2.3	NP languages . . . . .	20
2.3.1	Non-deterministic Turing Machines . . . . .	21
2.4	Reductions and hardness . . . . .	22
2.4.1	Example: Boolean formulas and the conjunctive normal form . . . . .	22
2.4.2	Example: reducing 3-SAT to INDSET . . . . .	24
<b>II</b>	<b>Questions and exercises</b>	<b>26</b>
<b>A</b>	<b>Self-assessment questions</b>	<b>27</b>
A.1	Computability . . . . .	27
A.1.1	Recursive and recursively enumerable sets . . . . .	27
A.1.2	Turing machines . . . . .	27
<b>B</b>	<b>Exercises</b>	<b>28</b>

# Part I

## Lecture notes

# Chapter 1

## Computability

### 1.1 Basic definitions and examples

In computer science, every problem instance can be represented by a finited sequence of symbols from a finite alphabet, or equivalently as a natural number. Therefore, we can restrict our focus on functions mapping natural numbers to natural numbers:

$$f : \mathbb{N} \rightarrow \mathbb{N}.$$

In particular, we are interested in functions for which the answer is “yes” or “no”, which can be modeled as

$$f : \mathbb{N} \rightarrow \{0, 1\};$$

in such case, we talk about a *decision problem*.

Examples:

- Given a natural number  $n$ , is  $n$  prime?
- Given a graph, what is the maximum degree of its nodes?
- From a customer database, select the customers that are more than fifty years old.
- Given a set of pieces of furniture and a set of trucks, can we accommodate all the furniture in the trucks?

As long as the function’s domain and codomain are finite, they can be represented as sequences of symbols, hence of bits, therefore as integer numbers (although some representations make more sense than others); observe that some problems among those listed are decision problems, others aren’t.

#### Decision functions and sets

There is a one-to-one correspondence between decision functions on the natural numbers and subsets of natural numbers. Given  $f : \mathbb{N} \rightarrow \{0, 1\}$ , its obvious set counterpart is the subset of natural numbers for which the function answers 1:

$$S_f = \{n \in \mathbb{N} : f(n) = 1\}.$$

Conversely, given a natural number subset  $S \subseteq \mathbb{N}$ , we can always define the function that decides over elements of the set:

$$f_S = \begin{cases} 1 & \text{if } n \in S \\ 0 & \text{if } n \notin S. \end{cases}$$

Given a function, or equivalently a set, we say that it is **computable**<sup>1</sup> (or **decidable**, or **recursive**) if and only if a procedure can be described to compute the function’s outcome in a finite number of steps. Observe that, in order for this definition to make sense, we need to define what an acceptable “procedure” is; for the time being, let us intuitively consider any computer algorithm.

Examples of computable functions and sets are the following:

- the set of even numbers;
- a function that decides whether a number is prime or not;
- any finite or cofinite<sup>2</sup> set, and any function that decides on them;
- any function studied in a basic Algorithms course (sorting, hashing, spanning trees on graphs. . .).

### 1.1.1 A possibly non-recursive set

#### Collatz numbers

Given  $n \in \mathbb{N} \setminus \{0\}$ , let the *Collatz sequence* starting from  $n$  be defined as follows:

$$a_1 = n$$

$$a_{i+1} = \begin{cases} a_i/2 & \text{if } a_i \text{ is even} \\ 3a_i + 1 & \text{if } a_i \text{ is odd,} \end{cases} \quad i = 1, 2, \dots$$

In other words, starting from  $n$ , we repeatedly halve it while it is even, and multiply it by 3 and add 1 if it is odd.

The *Collatz conjecture*<sup>3</sup> states that every Collatz sequence eventually reaches the value 1. While most mathematicians believe it to be true, nobody has been able to prove it.

Suppose that we are asked the following question:

“Given  $n \in \mathbb{N} \setminus \{0\}$ , does the Collatz sequence starting from  $n$  reach 1?”

If the answer is “yes,” let us call  $n$  a *Collatz number*. Let  $f : \mathbb{N} \setminus \{0\} \rightarrow \{0, 1\}$  be the corresponding decision function:

$$f(n) = \begin{cases} 1 & \text{if } n \text{ is a Collatz number} \\ 0 & \text{if } n \text{ is not a Collatz number,} \end{cases} \quad n = 1, 2, \dots$$

Then the Collatz conjecture simply states that all positive integers are Collatz numbers or, equivalently, that  $f(n) = 1$  on its whole domain.

#### Decidability of the Collatz property

Let us consider writing a function, in any programming language, to answer the above question, i.e., a function that returns 1 if and only if its argument is a Collatz number. Figure 1.1 details two possible ways to do it, and both have problems: the rightmost one requires us to have faith in an unproven mathematical conjecture; the left one only halts when the answer is 1 (the final **return** is never reached).

In more formal terms, we are admitting that we are **not** able to prove that the Collatz property is *decidable* (i.e., that there is a computer program that always terminates with the correct answer<sup>4</sup>). However, we have provided a procedure that terminates with the correct answer when the answer is

<sup>1</sup>[https://en.wikipedia.org/wiki/Recursive\\_set](https://en.wikipedia.org/wiki/Recursive_set)

<sup>2</sup>A set is *cofinite* when its complement is finite.

<sup>3</sup>[https://en.wikipedia.org/wiki/Collatz\\_conjecture](https://en.wikipedia.org/wiki/Collatz_conjecture)

<sup>4</sup>To the best of my knowledge, which isn’t much.

<pre> <b>function</b> collatz (<math>n \in \mathbb{N} \setminus \{0\}</math>) <math>\in \{0, 1\}</math>   <b>repeat</b>   [ <b>if</b> <math>n = 1</math> <b>then return</b> 1     [ <b>if</b> <math>n</math> is even       [ <b>then</b> <math>n \leftarrow n/2</math>         [ <b>else</b> <math>n \leftarrow 3n + 1</math>       ]     ]   ]   <b>return</b> 0 </pre>	<pre> <b>function</b> collatz (<math>n \in \mathbb{N} \setminus \{0\}</math>) <math>\in \{0, 1\}</math>   <b>return</b> 1 </pre>
--	--

Figure 1.1: Left: the only way I know to decide whether  $n$  is a Collatz number isn't guaranteed to work. Right: a much better way, but it is correct if and only if the conjecture is true.

“yes” (the function is not *total*, in the sense that it doesn't always provide an answer). We call such set **recursively enumerable**<sup>5</sup> (or RE, in short).

Having a procedure that only terminates when the answer is “yes” might not seem much, but it actually allows us to enumerate all numbers having the property. The function in Fig. 1.2 shows the basic trick to enumerate a potentially non-recursive set, applied to the Collatz sequence: the **diagonal method**<sup>6</sup>. Rather than performing the whole decision function on a number at a time (which would expose us to the risk of an endless loop), we start by executing the first step of the decision function for the first input ( $n = 1$ ), then we perform the second step for  $n = 1$  and the first step of  $n = 2$ ; at the  $i$ -th iteration, we perform the  $i$ -th step of the first input, the  $(i - 1)$ -th for the second, down to the first step for the  $i$ -th input. This way, every Collatz number will eventually hit 1 and be printed out.

The naïf approach of following the table rows is not guaranteed to work, since it would loop indefinitely, should a non-Collatz number ever exist.

Observe that the procedure does not print out the numbers in increasing order.

## 1.2 A computational model: the Turing machine

Among the many formal definition of computation proposed since the 1930s, the Turing Machine (TM for short) is by far the most similar to our intuitive notion. A Turing Machine<sup>7</sup> is defined by:

- a finite alphabet  $\Sigma$ , with a distinguished “default” symbol (e.g., “ $\_$ ” or “0”) whose symbols are to be read and written on an infinitely extended tape divided into cells;
- a finite set of states  $Q$ , with a distinguished initial state and one or more distinguished halting states;
- a set of rules  $R$ , described by a (possibly partial) function that associates to a pair of symbol and state a new pair of symbol and state plus a direction:

$$R : Q \times \Sigma \rightarrow \Sigma \times Q \times \{L, R\}.$$

Initially, all cells contain the default symbol, with the exception of a finite number; the non-blank portion of the tape represent the *input* of the TM. The machine also maintains a *current position* on the tape. The machine has an initial state  $q_0 \in Q$ . At every step, if the machine is in state  $q \in Q$ , and the symbol  $\sigma \in \Sigma$  appears in the current position of the tape, the machine applies the rule set  $R$  to  $(q, \sigma)$ :

$$(\sigma', q', d) = R(q, \sigma).$$

The machine writes the symbol  $\sigma'$  on the current tape cell, enters state  $q'$ , and moves the current position by one cell in direction  $d$ . If the machine enters one of the distinguished halting states, then

<sup>5</sup>[https://en.wikipedia.org/wiki/Recursively\\_enumerable\\_set](https://en.wikipedia.org/wiki/Recursively_enumerable_set)

<sup>6</sup>See <https://comp3.eu/collatz.py> for a Python version.

<sup>7</sup>[https://en.wikipedia.org/wiki/Turing\\_machine](https://en.wikipedia.org/wiki/Turing_machine)



```

1. procedure enumerate_collatz
2.   queue ← []
3.   for n ← 1 ... ∞
4.     queue_n ← n
5.     for i ← 1 ... n:
6.       if queue_i = 1
7.         print i
8.         delete queue_i
9.       else if queue_i is not deleted
10.        if queue_i is even
11.          then queue_i ← queue_i / 2
12.        else queue_i ← 3 · queue_i + 1

```

*Repeat for all numbers*  
*Add n to queue with itself as starting value*  
*Iterate on all numbers up to n*  
*i is Collatz, print and forget it*  
  
*deleted means "Already taken care of"*  
*if current number wasn't printed and forgotten yet*  
*Advance i-th sequence in the queue by one step*




Figure 1.2: Enumerating all Collatz numbers: top: the algorithm; bottom: a working schematic

the computation ends. At this point, the content of the (non-blank portion of) the tape represents the computation's *output*.

Observe that the input size for a TM is unambiguously defined: the size of the portion of tape that contains non-default symbols. Also the “execution time” is well understood: it is the number of steps before halting. Therefore, when we say that the computational complexity of a TM for inputs of size  $n$  is  $T(n)$  then we mean that  $T(n)$  is the worst-case number of steps that a TM performs before halting when the input has size  $n$ .

### 1.2.1 Examples

In order to experiment with Turing machines, many web-based simulators are available. The two top search results for “turing machine demo” are

- <http://morphett.info/turing/turing.html>
- <https://turingmachinesimulator.com/>.

Students are invited to read the simplest examples and to try implementing a TM for some simple problem (e.g., some arithmetic or logical operation on binary or unary numbers). Also, see the examples provided in the course web page.

### 1.2.2 Computational power of the Turing Machine

With reference to more standard computational models, such as the Von Neumann architecture of all modern computers, the TM seems very limited; for instance, it lacks any random-access capability. The next part of this course is precisely meant to convince ourselves that a TM is exactly as powerful as any other (theoretical) computational device. To this aim, let us discuss some possible generalizations.

#### Multiple-tape Turing machines

A  $k$ -tape Turing machine is a straightforward generalization of the basic model, with the following variations:

- the machine has  $k$  unlimited tapes, each with an independent current position;
- the rule set of the machine takes into account  $k$  symbols (one for each tape, from the current position) both in reading and in writing, and  $k$  movement directions (each current position is independent), with the additional provision of a “stay” direction  $S$  in which the position does not move:

$$R : Q \times \Sigma^k \rightarrow \Sigma^k \times Q \times \{L, R, S\}^k.$$

Multiple-tape TMs are obviously more practical for many problems. For example, try following the execution of the binary addition algorithms below:

- 1-tape addition from <http://morphett.info/turing/turing.html>: select “Load an example program/Binary addition”;
- 3-tape addition from <https://turingmachinesimulator.com/>: select “Examples/3 tapes/Binary addition”.

However, it turns out that any  $k$ -tape Turing machine can be “simulated” by a 1-tape TM, in the sense that it is possible to represent a  $k$ -tape TM on one tape, and to create a set of 1-tape rules that simulates the evolution of the  $k$ -tape TM. Of course, the 1-tape machine is much slower, as it needs to repeatedly scan its tape back and forth just to simulate a single step of the  $k$ -tape one.

**Theorem 1** (*k*-tape Turing machine emulation). *If a k-tape Turing machine  $\mathcal{M}$  takes time  $T(n)$  on inputs of time  $n$ , then it is possible to program a 1-tape Turing machine  $\mathcal{M}'$  that simulates it (i.e., essentially performs the same computation) in time  $O(T(n)^2)$ .*

*Proof.* See Arora-Barak, Claim 1.9 in the public draft.

Basically, the  $k$  tapes of  $\mathcal{M}$  are encoded on the single tape of  $\mathcal{M}'$  by alternating the cell contents of each tape; in order to remember the “current position” on each tape, every symbol is complemented by a different version (e.g., a “hatted” symbol) to be used to mark the current position. To emulate a step of  $\mathcal{M}$ , the whole tape of  $\mathcal{M}'$  is first scanned in order to find the  $k$  symbols in the current positions; then, a second scan is used to replace each symbol in the current position with the new symbol; then a third scan performs an update of the current positions.

Since  $\mathcal{M}$  halts in  $T(n)$  steps, no more than  $T(n)$  cells of the tapes will ever be visited; therefore, every scan performed by  $\mathcal{M}'$  will take at most  $kT(n)$  steps. Given some more details, cleanup tasks and so on, the simulation of a single step of  $\mathcal{M}$  will take at most  $5kT(n)$  steps by  $\mathcal{M}'$ , therefore the whole simulation takes  $5kT(n)^2$  steps. Since  $5k$  is constant wrt the input size  $n$ , the result follows.  $\square$

### Size of the alphabet

The number of symbols that can be written on a tape (the size of the alphabet  $\Sigma$ ) can make some tasks easier; for instance, in order to deal with binary numbers a three-symbol alphabet (“0”, “1”, and the blank as a separator) is convenient, while working on words is easier if the whole alphabet is available.

While a 1-sized alphabet  $\Sigma = \{\_ \}$  is clearly unfit for a TM (no way to store information on the tape), a 2-symbol alphabet is enough to simulate any TM:

**Theorem 2** (Emulation by a two-symbol Turing Machine). *If a Turing machine  $\mathcal{M}$  with a  $k$ -symbol alphabet  $\Sigma$  takes time  $T(n)$  on an input of size  $n$ , then it can be simulated by a Turing machine  $\mathcal{M}'$  with a 2-symbol alphabet  $\Sigma' = \{0, 1\}$  in time  $O(T(n))$  (i.e., with a linear slowdown).*

*Proof.* See Arora-Barak, claim 1.8 in the public draft, where for convenience machine  $\mathcal{M}'$  is assumed to have 4 symbols and the tape(s) extend only in one direction.

Every symbol from alphabet  $\Sigma$  can be encoded by  $\lceil \log_2 k \rceil$  binary digits from  $\Sigma'$ . Every step of machine  $\mathcal{M}$  will be simulated by  $\mathcal{M}'$  by reading  $\lceil \log_2 k \rceil$  cells in order to reconstruct the current symbol in  $\mathcal{M}$ ; the symbol being reconstructed bit by bit is stored in the machine state (therefore,  $\mathcal{M}'$  requires many more states than  $\mathcal{M}$ ). This scan is followed by a new scan to replace the encoding with the new symbol (again, all information needed by  $\mathcal{M}'$  will be “stored” in its state), and a third (possibly longer) scan to place the current position to the left or right encoding. Therefore, a step of  $\mathcal{M}$  will require less than  $4\lceil \log_2 k \rceil$  steps of  $\mathcal{M}'$ , and the total simulation time will be

$$T'(n) \leq 4\lceil \log_2 k \rceil T(n).$$

$\square$

### 1.2.3 Universal Turing machines

The main drawback of TMs, as described up to now, with respect to our modern understanding of computational systems, is that each serves one specific purpose, encoded in its rule set: a machine to add numbers, one to multiply, and so on.

However, it is easy to see that a TM can be represented by a finite string in a finite alphabet: each transition rule can be seen as a quintuplet, each from a finite set, and the set of rules is finite. Therefore, it is possible to envision a TM  $\mathcal{U}$  that takes another TM  $\mathcal{M}$  as input on its tape, properly encoded, together with an input string  $s$  for  $\mathcal{M}$ , and simulates  $\mathcal{M}$  step by step on input  $s$ . Such machine is called a Universal Turing machine (UTM).

One such machine, using a 16 symbol encoding and a single tape, is described in

[https://www.dropbox.com/sh/u7jsxm232giwown/AADTRNqjKBIE\\_QZGyicoZWjYa/utm.pdf](https://www.dropbox.com/sh/u7jsxm232giwown/AADTRNqjKBIE_QZGyicoZWjYa/utm.pdf)

and can be seen in action at the aforementioned link <http://morphett.info/turing/turing.html>, clicking “Load an example program / Universal Turing machine.”

### Simulating other computational devices

Now we know that TMs are powerful enough to, in a sense, “simulate themselves.” We can also convince ourselves quite easily that they can emulate a simple CPU/RAM architecture: just replace random access memory with sequential search on a tape (tremendous slowdown, but we are not concerned by it now), the CPU’s internal registers can be stored in separate tapes, and every opcode of the CPU corresponds to a separate set of states of the machine. Operations such as “load memory to a register,” “perform an arithmetic or logical operation between registers,” “conditionally jump to memory” and so on can be emulated.

### 1.2.4 The Church-Turing thesis

We should be convinced, by now, that TMs are powerful enough to be a fair computational model, at least on par with any other reasonable definition. We formalize this idea into a sort of “postulate,” i.e., an assertion that we will assume to be true for the rest of this course.

**Postulate 1** (Church-Turing thesis). *Turing machines are at least as powerful as every physically realizable model of computation.*

This thesis allows us to extend every the validity negative result about TMs to every physical computational device.

## 1.3 Uncomputable functions

It is easy to understand that, even if we restrict our interest to decision functions, almost all functions are not computable by a TM. In fact, as the following Lemmata 1 and 2 show, there are simply too many functions to be able to define a TM for each of them.

**Lemma 1.** *The set of decision functions  $f : \mathbb{N} \rightarrow \{0,1\}$  (or, equivalently,  $f : \Sigma^* \rightarrow \{0,1\}$ ), is uncountable.*

*Proof.* By contradiction, suppose that a complete mapping exists from the naturals to the set of decision functions; i.e., there is a mapping  $n \mapsto f_n$  that enumerates all functions. Define function  $\hat{f}(n) = 1 - f_n(n)$ . By definition, function  $\hat{f}$  differs from  $f_n$  on the value it is assigned for  $n$  (if  $f_n(n) = 0$  then  $\hat{f}(n) = 1 - f_n(n) = 1 - 0 = 1$ , and vice versa). Therefore, contrary to the assumption, the enumeration is not complete because it excluded function  $\hat{f}$ .  $\square$

Lemma 1 is an example of *diagonal argument*, introduced by Cantor in order to prove the uncountability of real numbers: focus on the “diagonal” values (in our case  $f_n(n)$ , by using the same number as function index and as argument), and make a new object that systematically differs from all that are listed.

**Lemma 2.** *Given a finite alphabet  $\Sigma$ , the number of TMs on that alphabet is countable.*

*Proof.* As shown in the Universal TM discussion, every TM can be encoded in some appropriate alphabet. As shown by Theorem 2, every alphabet with at least two symbols can emulate and be emulated by every other alphabet. Therefore, it is possible to define a representation of any TM in any alphabet.

We know that strings can be enumerated: first we count the only string in  $\Sigma^0$ , then the strings in  $\Sigma^1$ , then those in  $\Sigma^2$  (e.g., in lexicographic order), and so on. Since every string  $s \in \Sigma^*$  is finite

( $s \in \Sigma^{|s|}$ ), sooner or later it will be enumerated. Therefore there is a mapping  $\mathbb{N} \rightarrow \Sigma^*$ , i.e.,  $\Sigma^*$  is countable.

Since TMs can be mapped on a subset of  $\Sigma^*$  (those strings that define TMs according to the chosen encoding), and are still infinite, it follows that TMs are countable.  $\square$

Therefore, whatever way we choose to enumerate TMs and to associate them with decision functions, we will inevitably leave out some functions. Hence, given that TMs are our definition of computing,

**Corollary 1.** *There are uncomputable decision functions.*

### 1.3.1 Finding an uncomputable function

Let us introduce a little more notation. As already defined, the alphabet  $\Sigma$  contains a distinguished, “default” symbol, which we assume to be “\_”. Before the computation starts, only a finite number of cell tapes have non-blank symbols. Let us define as “input” the smallest, contiguous set of tape cells that contains all non-blank symbols.

A Turing machine transforms an input string into an output string (the smallest contiguous set of tape cells that contain all non-blank symbols at the *end* of the computation), but it might never terminate. In other words, if we see a TM machine as a function from  $\Sigma^*$  to  $\Sigma^*$  it might not be a *total* function.

As an alternative, we may introduce a new value,  $\infty$ , as the “value” of a non-terminating computation; given a Turing machine  $\mathcal{M}$ , if its computation on input  $s$  does not terminate we will write  $\mathcal{M}(s) = \infty$ .

While TM encodings have a precise syntax, so that not all strings in  $\Sigma^*$  are syntactically valid encodings of some TM, we can just accept the convention that any such invalid string encodes the TM that immediately halts (think of  $s$  as a program, executed by a UTM that immediately stops if there is a syntax error). This way, all strings can be seen to encode a TM, and most string just encode the “identity function” (a machine that halts immediately leaves its input string unchanged). Let us therefore call  $\mathcal{M}_s$  the TM whose encoding is string  $s$ , or the machine that immediately terminates if  $s$  is not a valid encoding.

With this convention in mind, we can design a function whose outcome differs from that of any TM. We employ a diagonal technique akin to the proof of Lemma 1: for any string  $\alpha \in \Sigma^*$ , we define our function to differ from the output of the TM encoded by  $\alpha$  on input  $\alpha$  itself.

**Theorem 3.** *Given an alphabet  $\Sigma$  and a encoding  $\alpha \mapsto \mathcal{M}_\alpha$  of TMs in that alphabet, the function*

$$UC(\alpha) = \begin{cases} 0 & \text{if } \mathcal{M}_\alpha(\alpha) = 1 \\ 1 & \text{otherwise} \end{cases} \quad \forall \alpha \in \Sigma^*$$

*is uncomputable.*

*Proof.* Let  $\mathcal{M}$  be any TM, and let  $m \in \Sigma^*$  be its encoding (i.e.,  $\mathcal{M} = \mathcal{M}_m$ ). By definition,  $UC(m)$  differs from  $\mathcal{M}(m)$ : the former outputs one if and only if the latter outputs anything else (or does not terminate).

See also Arora-Barak, theorem 1.16 in the public draft.  $\square$

What is the problem that prevents us from computing  $UC$ ? While the definition is quite straightforward, being able to emulate the machine  $\mathcal{M}_\alpha$  on input  $\alpha$  is not enough to always decide the value of  $UC(\alpha)$ . We need to take into account also the fact that the emulation might never terminate. This allows us to prove, as a corollary of the preceding theorem, that there is no procedure that always determines whether a machine will terminate on a given input.

**Theorem 4** (Halting problem). *Given an alphabet  $\Sigma$  and a encoding  $\alpha \mapsto \mathcal{M}_\alpha$  of TMs in that alphabet, the function*

$$HALT(s, t) = \begin{cases} 0 & \text{if } \mathcal{M}_s(t) = \infty \\ 1 & \text{otherwise} \end{cases} \quad \forall (s, t) \in \Sigma^* \times \Sigma^*$$

(i.e., which returns 1 if and only if machine  $\mathcal{M}_s$  halts on input  $t$ ) is uncomputable.

*Proof.* Let's proceed by contradiction. Suppose that we have a machine  $\mathcal{H}$  which computes  $HALT(s, t)$  (i.e., when run on a tape containing a string  $s$  encoding a TM and a string  $t$ , always halts returning 1 if machine  $\mathcal{M}_s$  would halt when run on input  $t$ , and returning 0 otherwise). Then we could use  $\mathcal{H}$  to compute function  $UC$ .

For convenience, let us compute  $UC$  using a machine with two tapes. The first tape is read-only and contains the input string  $\alpha \in \Sigma^*$ , while the second will be used as a work (and output) tape. To compute  $UC$ , the machine will perform the following steps:

- Create two copies of the input string  $\alpha$  onto the work tape, separated by a blank (we know we can do this because we can actually write the machine);
- Execute the machine  $\mathcal{H}$  (which exists by hypothesis) on the work tape, therefore calculating whether the computation  $\mathcal{M}_\alpha(\alpha)$  would terminate or not. Two outcomes are possible:
  - If the output of  $\mathcal{H}$  is zero, then we know that the computation of  $\mathcal{M}_\alpha(\alpha)$  wouldn't terminate, therefore, by definition of function  $UC$ , we can output 1 and terminate.
  - If, on the other hand, the output of  $\mathcal{H}$  is one, then we know for sure that the computation  $\mathcal{M}_\alpha(\alpha)$  would terminate, and we can emulate it with a UTM  $\mathcal{U}$  (which we know to exist) and then “inverting” the result à la  $UC$ , by executing the following steps:
    - \* As in the first step, create two copies of the input string  $\alpha$  onto the work tape, separated by a blank;
    - \* Execute the UTM  $\mathcal{U}$  on the work tape, thereby emulating the computation  $\mathcal{M}_\alpha(\alpha)$ ;
    - \* At the end, if the output of the emulation was 1, then replace it by a 0; if it was anything other than 1, replace it with 1.

This machine would be able to compute  $UC$  by simply applying its definition, but we know that  $UC$  is not computable by a TM; all steps, apart from  $\mathcal{H}$ , are already known and computable. We must conclude that  $\mathcal{H}$  cannot exist.

See also Arora-Barak, theorem 1.17 in the public draft. □

This proof employs a very common technique of CS, called *reduction*: in order to prove the impossibility of  $HALT$ , we “reduce” the computation of  $UC$  to that of  $HALT$ ; since we know that the former is impossible, we must conclude that the latter is too.

### The Halting Problem for machines without an input

Consider the special case of machines that do not work on an input string; i.e., the class of TMs that are executed on a completely blank tape. Asking whether a computation without input will eventually halt might seem a simpler question, because we somehow restrict the number of machines that we are considering.

Let us define the following specialized halting function:

$$HALT'(s) = HALT(s, \varepsilon) = \begin{cases} 0 & \text{if } \mathcal{M}_s(\varepsilon) = \infty \\ 1 & \text{otherwise} \end{cases} \quad \forall s \in \Sigma^*.$$

It turns out that if we were able to compute  $HALT'$  then we could also compute  $HALT$ :

**Theorem 5.** *HALT' is not computable.*

*Proof.* By contradiction, suppose that there is a machine  $\mathcal{H}'$  that computes  $HALT'$ . Such machine would be executed on a string  $s$  on the tape, and would return 1 if the machine encoded by  $s$  would halt when run on an empty tape, 0 otherwise.

Now, suppose that we are asked to compute  $HALT(s, t)$  for a non-empty input string  $t$ . We can transform the computation  $\mathcal{M}_s(t)$  on a computation  $\mathcal{M}_{s'}(\varepsilon)$  on an empty tape where  $s'$  contains the whole encoding  $s$ , but prepended with a number of states that write the string  $t$  on the tape. In other words, we transform a computation on a generic input into a computation on an empty tape that writes the desired input before proceeding.

After modifying the string  $s$  into  $s'$  on tape, we can run  $\mathcal{H}'$  on it. The answer of  $\mathcal{H}'$  is precisely  $HALT(s, t)$ , which would therefore be computable.  $\square$

Again, the result was obtained by reducing a known impossible problem,  $HALT$  to the newly introduced one,  $HALT'$ .

### Consequences of the Halting Problem incomputability

If  $HALT$  were computable, we would be able to settle any mathematical question that can be disproved by a counterexample (on a discrete set), such as the Collatz conjecture, Goldbach's conjecture<sup>8</sup>, the non-existence of odd perfect numbers<sup>9</sup>... We would just need to write a machine that systematically search for one such counterexample and halts as soon as it finds one: by feeding this machine as an input to  $\mathcal{H}$ , we would know whether a counterexample exists at all or not.

More generally, for every proposition  $P$  in Mathematical logic we would know whether it is provable or not: just define a machine that, starting from pre-encoded axioms, systematically generates all their consequences (theorems) and halts whenever it generates  $P$ . Machine  $\mathcal{H}$  would tell us whether  $P$  is ever going to be generated or not.

Note that, in all cases described above, we would only receive a "yes/no" answer, not an actual counterexample or a proof.

### 1.3.2 Recursive enumerability of halting computations

Although  $HALT$  is not computable, it is clearly recursively enumerable. In fact, we can just take a UTM and modify it to erase the tape and write "1" whenever the emulated machine ends, and we would have a partial function that always accepts (i.e., returns 1) on terminating computations.

It is also possible to output all  $(s, t) \in \Sigma^* \times \Sigma^*$  pairs for which  $\mathcal{M}_s(t)$  halts by employing a diagonal method similar to the one used in Fig. 1.2<sup>10</sup>.

Function  $HALT$  is our first example of R.E. function that is provably not recursive.

Observe that, unlike recursivity, R.E. does *not* treat the "yes" and "no" answer in a symmetric way. We can give the following:

**Definition 1.** *A decision function  $f : \Sigma^* \rightarrow \{0, 1\}$  is co-R.E. if it admits a TM  $\mathcal{M}$  such that  $\mathcal{M}(s)$  halts with output 0 if and only if  $f(s) = 0$ .*

In other words, co-R.E. functions are those for which it is possible to compute a "no" answer, while the computation might not terminate if the answer is "yes". Clearly, if  $f$  is R.E., then  $1 - f$  is co-R.E.

**Theorem 6.** *A decision function  $f : \Sigma^* \rightarrow \{0, 1\}$  is recursive if and only if it is both R.E. and co-R.E.*

<sup>8</sup>Every even number (larger than 2) can be expressed as the sum of two primes, see [https://en.wikipedia.org/wiki/Goldbach%27s\\_conjecture](https://en.wikipedia.org/wiki/Goldbach%27s_conjecture)

<sup>9</sup>[https://en.wikipedia.org/wiki/Perfect\\_number](https://en.wikipedia.org/wiki/Perfect_number)

<sup>10</sup>See the figure at [https://en.wikipedia.org/wiki/Recursively\\_enumerable\\_set#Examples](https://en.wikipedia.org/wiki/Recursively_enumerable_set#Examples)

*Proof.* Let us prove the “only if” part first. If  $f$  is recursive, then there is a TM  $\mathcal{M}_f$  that computes it. But  $\mathcal{M}_f$  clearly satisfies both the R.E. definition ( $\mathcal{M}_f(s)$  halts with output 1 if and only if  $f(s) = 1$ ) and the co-R.E. definition ( $\mathcal{M}_f(s)$  halts with output 0 if and only if  $f(s) = 0$ ).

About the “if” part, if  $f$  is R.E., then there is a TM  $M_1$  such that  $M_1(s)$  halts with output 1 iff  $f(s) = 1$ ; since  $f$  is also co-R.E., then there is also a TM  $\mathcal{M}_0$  such that  $M_1(s)$  halts with output 0 iff  $f(s) = 0$ . Therefore, a machine that alternates one step of the execution of  $\mathcal{M}_1$  with one step of  $\mathcal{M}_0$ , halting when one of the two machines halts and returning its output, will eventually terminate (because, whatever the value of  $f$ , at least one of the two machines is going to eventually halt) and precisely decides  $f$ .  $\square$

Observe that, as already pointed out, any definition given on decision functions with domain  $\Sigma^*$  also works on domain  $\mathbb{N}$  (and on any other discrete domain), and can be naturally extended on subsets of strings or natural numbers. We can therefore define a set as recursive, recursively enumerable, or co-recursively enumerable.

### 1.3.3 Another uncomputable function: the Busy Beaver game

Since we might be unable to tell at all whether a specific TM will halt, the question arises of how complex can machine’s output be for a given number of states.

**Definition 2** (The Busy Beaver game). *Among all TMs on alphabet  $\{0, 1\}$  and with  $n = |Q|$  states (not counting the halting one) that halt when run on an empty (i.e., all-zero) tape:*

- let  $\Sigma(n)$  be the largest number of (not necessarily consecutive) ones left by any machine upon halting;
- let  $S(n)$  be the largest number of steps performed by any such machine before halting.

Function  $\Sigma(n)$  is known as the *busy beaver* function for  $n$  states, and the machine that achieves it is called the Busy Beaver for  $n$  states.

Both functions grow very rapidly with  $n$ , and their values are only known for  $n \leq 4$ . The current Busy Beaver candidate with  $n = 5$  states writes more than 4K ones before halting after more than 47M steps.

**Theorem 7.** *The function  $S(n)$  is not computable.*

*Proof.* Suppose that  $S(n)$  is computable. Then, we could create a TM to compute  $HALT'$  (the variant with empty input) on a machine encoded in string  $s$  as follows:

- Count the number  $n$  of states of  $\mathcal{M}_s$ ;
- compute  $\ell = S(n)$ ;
- emulate  $\mathcal{M}_s$  for  $\ell$  steps;
  - if the emulation halts before, then  $\mathcal{M}_s$  clearly halts: write 1 on the tape and halt;
  - otherwise, since  $\mathcal{M}_s$  has been running longer than the longest-running machine with  $n$  states, it is bound to run forever: write 0 on the tape and halt.

$\square$

The next result is even stronger. Given two functions  $f, g : \mathbb{N} \rightarrow \mathbb{N}$ , we say that  $f$  “eventually outgrows”  $g$ , written  $f >_E g$ , if  $f(n) \geq g(n)$  for a sufficiently large value of  $n$ :

$$f >_E g \Leftrightarrow \exists N : \forall n > N f(n) \geq g(n).$$

**Theorem 8.** *The function  $\Sigma(n)$  eventually outgrows any computable function.*



*Proof.* Let  $f : \mathbb{N} \rightarrow \mathbb{N}$  be computable. Let us define the following function:

$$F(n) = \sum_{i=0}^n [f(i) + i^2].$$

By definition,  $F$  clearly has the following properties:

$$F(n) \geq f(n) \quad \forall n \in \mathbb{N}, \tag{1.1}$$

$$F(n) \geq n^2 \quad \forall n \in \mathbb{N}, \tag{1.2}$$

$$F(n+1) > F(n) \quad \forall n \in \mathbb{N} \tag{1.3}$$

the latter because  $F(n+1)$  is equal to  $F(n)$  plus a strictly positive term. Moreover, since  $f$  is computable,  $F$  is computable too. Suppose that  $M_F$  is a TM on alphabet  $\{0, 1\}$  that, when positioned on the rightmost symbol of an input string of  $x$  ones and executed, outputs a string of  $F(x)$  ones (i.e., computes the function  $x \mapsto F(x)$  in unary representation) and halts below the rightmost one. Let  $C$  be the number of states of  $M_F$ .

Given an arbitrary integer  $x \in \mathbb{N}$ , we can define the following machine  $\mathcal{M}$  running on an initially empty tape (i.e., a tape filled with zeroes):

- Write  $x$  ones on the tape and stop at the rightmost one (i.e., the unary representation of  $x$ : it can be done with  $x$  states, see Exercise 3 at page 31);
- Execute  $M_F$  on the tape (therefore computing  $F(x)$  with  $C$  states);
- Execute  $M_F$  again on the tape (therefore computing  $F(F(x))$  with  $C$  more states).

The machine  $\mathcal{M}$  works on alphabet  $\{0, 1\}$ , starts with an empty tape, ends with  $F(F(x))$  ones written on it and has  $x + 2C$  states; therefore it is a busy beaver candidate, and the  $(x + 2C)$ -state busy beaver must perform at least as well:

$$\Sigma(x + 2C) \geq F(F(x)). \tag{1.4}$$

Now,

$$F(x) \geq x^2 >_E x + 2C;$$

the first inequality comes from (1.2), while the second stems from the fact that  $x^2$  eventually dominates any linear function of  $x$ . By applying  $F$  to both the left- and right-hand sides, which preserves the inequality sign because of (1.3), we get

$$F(F(x)) >_E F(x + 2C). \tag{1.5}$$

By concatenating (1.4), (1.5) and (1.1), we get

$$\Sigma(x + 2C) \geq F(F(x)) >_E F(x + 2C) \geq f(x + 2C).$$

Finally, by replaxing  $n = x + 2C$ , we obtain

$$\Sigma(n) >_E f(n).$$

□

This proof is based on the original one by Rado (1962)<sup>11</sup>.

---

<sup>11</sup>See for instance:  
[http://computation4cognitivescientists.weebly.com/uploads/6/2/8/3/6283774/rado-on\\_non-computable\\_functions.pdf](http://computation4cognitivescientists.weebly.com/uploads/6/2/8/3/6283774/rado-on_non-computable_functions.pdf)

### 1.3.4 More definitions

A few more definitions, for brevity:

- Given an alphabet  $\Sigma$ , a set of strings  $S \subseteq \Sigma^*$  is also called a **language**.
- If language  $S$  is recursively enumerable, i.e. there is a TM  $\mathcal{M}$  such that  $\mathcal{M}(s) = 1 \Leftrightarrow s \in S$ , then we say that  $\mathcal{M}$  **accepts**  $S$  (or that it “recognizes” it).
- Given a TM  $\mathcal{M}$ , the language recognized by it (i.e., the set of all inputs that are accepted by the machine) is represented by  $L(\mathcal{M})$ .
- If language  $S$  is recursive, i.e. there is a TM  $\mathcal{M}$  that accepts it and always halts, then we say that  $\mathcal{M}$  **decides**  $S$ .

### 1.3.5 Rice’s Theorem

Among all questions that we may ask about a Turing machine  $\mathcal{M}$ , some of them have a *syntactic* nature, i.e., they refer to its actual implementation: “does  $\mathcal{M}$  halt within 50 steps?”, “Does  $\mathcal{M}$  ever reach state  $q$ ?”, “Does  $\mathcal{M}$  ever print symbol  $\sigma$  on the tape?”...

Other questions are of a *semantic* type, i.e., they refer to the language accepted by  $\mathcal{M}$ , with no regards about  $\mathcal{M}$ ’s behavior: “does  $\mathcal{M}$  only accept even-length strings?”, “Does  $\mathcal{M}$  accept any string?”, “Does  $\mathcal{M}$  accept at least 100 different strings?”...

**Definition 3.** A property of a TM is a mapping  $P$  from TMs to  $\{0, 1\}$ , and we say that  $\mathcal{M}$  has property  $P$  when  $P(\mathcal{M}) = 1$ .

**Definition 4.** A property is semantic if its value is shared by all TMs recognizing the same language: if  $L(\mathcal{M}) = L(\mathcal{M}')$ , then  $P(\mathcal{M}) = P(\mathcal{M}')$ .

By extension, we can say that a language  $S$  has a property  $P$  if the machine that recognizes  $S$  has. Finally, we define a property as *trivial* if all TMs have it, or if no TM has it. A property is *non-trivial* if there is at least one machine having it, and one not having it.

The two trivial properties (the one possessed by all TMs and the one possessed by none) are easy to decide, respectively by the machine that always accepts and by the one that always rejects. On the other hand:

**Theorem 9** (Rice’s Theorem). All non-trivial semantic properties of TMs are undecidable.

*Proof.* As usual, let’s work by contradiction via reduction from the Halting Problem.

Suppose that a non-trivial semantic property  $P$  is decidable; this means that there is a TM  $\mathcal{M}_P$  that can be run on the encoding of any TM  $\mathcal{M}$  and returns 1 if  $\mathcal{M}$  has property  $P$ , 0 otherwise.

Let us also assume that the empty language  $\emptyset$  does not have the property  $P$  (otherwise we can work on the complementary property), and that the Turing machine  $\mathcal{N}$  has the property  $P$  (we can always find  $\mathcal{N}$  because  $P$  is not trivial).

Given the strings  $s, t \in \Sigma^*$ , we can then check whether  $M_s(t)$  halts by building the following auxiliary TM  $\mathcal{N}'$  that, on input  $u$ , works as follows:

- move the input  $y$  onto an auxiliary tape for later use, and replace it with  $t$ ;
- execute  $M_s$  on input  $t$ ;
- when the simulation halts (which, as we know, might not happen), restore the original input  $u$  on the tape by copying it back from the auxiliary tape;
- run  $\mathcal{N}$  on the original input  $u$ .

The machine  $\mathcal{N}'$  we just defined accepts the same language as  $\mathcal{N}$  if  $M_s(t)$  halts, otherwise it runs forever, therefore accepting the empty language. Therefore, running our hypothetical decision procedure  $\mathcal{M}_P$  on machine  $\mathcal{N}'$  we obtain “yes” if  $M_s(t)$  halts (since in this case  $L(\mathcal{N}) = L(\mathcal{N}')$ ), and “no” if  $M_s(t)$  doesn't halt (and thus the empty language, which doesn't have the property  $P$ , is recognized).  $\square$

Observe that we simply use  $\mathcal{N}$ , which has the property, as a sort of Trojan horse for computation  $M_s(t)$ . See also the Wikipedia entry for Rice's Theorem<sup>12</sup>.

---

<sup>12</sup>[https://en.wikipedia.org/wiki/Rice%27s\\_theorem](https://en.wikipedia.org/wiki/Rice%27s_theorem)

## Chapter 2

# Complexity classes: P and NP

From now on, we will be only dealing with computable functions; the algorithms that we will analyze will always terminate, and our main concern will be about the amount of resources (time, space) required to compute them.

### 2.1 Definitions

When discussing complexity, we are mainly interested in the relationship between the size of the input and the execution “time” of an algorithm executed by a Turing machine. We still refer to TMs because both input size and execution time can be defined unambiguously in that model.

#### Input size

By “size” of the input, we mean the number of symbols used to encode it in the machine’s tape. Since we are only concerned in asymptotic relationships, the particular alphabet used by a machine is of no concern, and we may as well just consider machines with alphabet  $\Sigma = \{0, 1\}$ .

We require that the input data are encoded in a reasonable way. For instance, numbers may be represented in base-2 notation (although the precise base does not matter when doing asymptotic analysis), so that the size of the representation  $r_2(n)$  of integer  $n$  in base 2 is logarithmic with respect to its value:

$$|r_2(n)| = O(\log n).$$

In this sense, unary representations (representing  $n$  by a string of  $n$  consecutive 1’s) is not to be considered reasonable because its size is exponential with respect to the base-2 notation.

#### Execution time

We dub “execution time,” or simply “time,” the number of steps required by a TM to get to a halting state. Let  $\mathcal{M}$  be a TM that always halts. We can define the “time” function

$$\begin{aligned} t_{\mathcal{M}} : \Sigma^* &\rightarrow \mathbb{N} \\ x &\mapsto \# \text{ of steps before } \mathcal{M} \text{ halts on input } x \end{aligned}$$

that maps every input string  $x$  onto the number of steps that  $\mathcal{M}$  performs upon input  $x$  before halting.  $\mathcal{M}$  always halts, so it is a well-defined function. Since the number of strings of a given size  $n$  is finite, we can also define (and actually compute, if needed) the following “worst-case” time for inputs of size  $n$ :

$$\begin{aligned} T_{\mathcal{M}} : \mathbb{N} &\rightarrow \mathbb{N} \\ n &\mapsto \max\{t_{\mathcal{M}}(x) : x \in \Sigma^n\}, \end{aligned}$$

i.e.,  $T_{\mathcal{M}}(n)$  is the longest time that  $\mathcal{M}$  takes before halting on an input of size  $n$ .

## 2.2 Polynomial languages

Let us now focus on decision problems.

**Definition 5.** Let  $f : \mathbb{N} \rightarrow \mathbb{N}$  be any computable function. We say that a language  $L \subseteq \Sigma^*$  is of class  $DTIME(f)$ , and write  $L \in DTIME(f)$ , if there is a TM  $\mathcal{M}$  that decides  $L$  and its worst-case time, as a function of input size, is dominated by  $f$ :

$$L \in DTIME(f) \iff \exists \mathcal{M} : L(\mathcal{M}) = L \wedge T_{\mathcal{M}} = O(f).$$

In other words,  $DTIME(f)$  is the class of all languages that can be decided by some TM in time eventually bounded by function  $c \cdot f$ , where  $c$  is constant.

Saying  $L \in DTIME(f)$  means that there is a machine  $\mathcal{M}$ , a constant  $c \in \mathbb{N}$  and an input size  $n_0 \in \mathbb{N}$  such that, for every input  $x$  with size larger than  $n_0$ ,  $\mathcal{M}$  decides  $x \in L$  in at most  $c \cdot f(|x|)$  steps.

Languages that can be decided in a time that is polynomial with respect to the input size are very important, so we give a short name to their class:

**Definition 6.**

$$\mathbf{P} = \bigcup_{k=0}^{\infty} DTIME(n^k).$$

In other words, we say that a language  $L \in \Sigma^*$  is polynomial-time, and write  $L \in \mathbf{P}$ , if there are a machine  $\mathcal{M}$  and a polynomial  $p(n)$  such that for every input string  $x$

$$x \in L \iff \mathcal{M}(x) = 1 \wedge t_{\mathcal{M}}(x) \leq p(|x|). \tag{2.1}$$

### 2.2.1 Examples

Here are some examples of polynomial-time languages.

**CONNECTED** — Given an encoding of graph  $G$  (e.g., the number of nodes followed by an adjacency matrix or list),  $G \in \text{CONNECTED}$  if and only if there is a path in  $G$  between every pair of nodes.

**PRIME** — Given a base-2 representation of a natural number  $N$ , we say that  $N \in \text{PRIME}$  if and only if  $N$  is, of course, prime.

Observe that the naive algorithm “divide by all integers from 2 to  $\lfloor \sqrt{N} \rfloor$ ” is *not* polynomial with respect to the size of the input string. In fact, the input size is  $n = O(\log N)$  (the number of bits used to represent a number is logarithmic with respect to its magnitude), therefore the naive algorithm would take  $\lfloor \sqrt{N} \rfloor - 1 = O(2^{n/2})$  divisions in the worst case, which grows faster than any polynomial<sup>1</sup>.

Anyway, it has recently been shown<sup>2</sup> that  $\text{PRIME} \in \mathbf{P}$ .

#### (Counter?)-examples

On the other hand, we do not know of any polynomial-time algorithm for the following languages:

**SATISFIABILITY or SAT** — (see also Sec. 2.4.1) Given a Boolean expression  $f(x_1, \dots, x_n)$  (usually in conjunctive normal form, CNF<sup>3</sup>) involving  $n$  variables, is there a truth assignment to the variables that satisfies (i.e., makes true) the formula<sup>4</sup>?

<sup>1</sup>An algorithm that is polynomial with respect to the *magnitude* of the numbers instead than the size of their representation is said to be “pseudo-polynomial.” In fact, the naive primality test would be polynomial if we chose to represent  $N$  in unary notation ( $N$  consecutive 1’s).

<sup>2</sup>[https://en.wikipedia.org/wiki/Primality\\_test#Fast\\_deterministic\\_tests](https://en.wikipedia.org/wiki/Primality_test#Fast_deterministic_tests)

<sup>3</sup>[https://en.wikipedia.org/wiki/Conjunctive\\_normal\\_form](https://en.wikipedia.org/wiki/Conjunctive_normal_form)

<sup>4</sup>[https://en.wikipedia.org/wiki/Boolean\\_satisfiability\\_problem](https://en.wikipedia.org/wiki/Boolean_satisfiability_problem)

**INDEPENDENT SET or INDSET** — Given an encoding of graph  $G$  and a number  $k$ , does  $G$  contain  $k$  nodes that are not connected to each other<sup>5</sup>?

**TRAVELING SALESMAN PROBLEM or TSP** — Given an encoding of a complete *weighted* graph  $G$  (i.e., all pairs of nodes are connected, and pair  $i, j$  is assigned a “weight”  $w_{ij}$ ) and a “budget”  $k$ , is there an order of visit (permutation)  $\sigma$  of all nodes such that

$$\left( \sum_{i=1}^{n-1} w_{\sigma_i \sigma_{i+1}} \right) + w_{\sigma_n \sigma_1} \leq k, \quad (2.2)$$

i.e., the total weight along the corresponding closed path in that order of visit (also considering return to the starting node) is within budget<sup>6</sup>?

However, we have no proof that these languages (and many others) are not in **P**. In the following section, we will try to characterize these languages.

## 2.3 NP languages

While the three languages listed above (SAT, INDSET, TSP) cannot be decided by any known polynomial algorithm, they share a common property: if a string is in the language, there is an “easily” (polynomially) verifiable proof of it:

- If  $f(x_1, \dots, x_n) \in \text{SAT}$  (i.e., boolean formula  $f$  is satisfiable), then there is a truth assignment to the variables  $x_1, \dots, x_n$  that satisfies it. If we were given this truth assignment, we could easily check that, indeed,  $f \in \text{SAT}$ . Note that the truth assignment consists of  $n$  truth values (bits) and is therefore shorter than the encoding of  $f$  (which contains a whole boolean expression on  $n$  variables), and that computing a Boolean formula can be reduced to a finite number of scans.
- If  $G \in \text{INDSET}$ , then there is a list of  $k$  independent nodes; given that list, we could easily verify that  $G$  does not contain any edge between them. The list contains  $k$  integers from 1 to the number of nodes in  $G$  (which is polynomial with respect to the size of  $G$ ’s representation) and requires a presumably quadratic or cubic time to be checked.
- If  $G \in \text{TSP}$ , then there is a permutation of the nodes in  $G$ , i.e., a list of nodes. Given that list, we can easily sum the weights as in (2.2) and check that the inequality holds.

In other words, if we are provided a *certificate* (or *witness*), it is easy for us to check that a given string belongs to the language. What’s important is that both the certificate’s size and the time to check are polynomial with respect to the input size. The class of such problems is called **NP**. More formally:

**Definition 7.** We say that a language  $L \subseteq \Sigma^*$  is of class **NP**, and write  $L \in \text{NP}$ , if there is a TM  $\mathcal{M}$  and two polynomials  $p(n)$  and  $q(n)$  such that for every input string  $x$

$$x \in L \iff \exists c \in \Sigma^{q(|x|)} : \mathcal{M}(x, c) = 1 \wedge t_{\mathcal{M}}(x, u) \leq p(|x|). \quad (2.3)$$

Basically, the two polynomials are needed to bound both the size of certificate  $c$  and the execution time of  $\mathcal{M}$ .

Observe that the definition only requires a (polynomially verifiable) certificate to exist only for “yes” answers, while “no” instances (i.e., strings  $x$  such that  $x \notin L$ ) might not be verifiable.

**Theorem 10.**

$$P \subseteq NP \subseteq \bigcup_{k=1}^{\infty} DTIME(2^{n^k}).$$

<sup>5</sup>[https://en.wikipedia.org/wiki/Independent\\_set\\_\(graph\\_theory\)](https://en.wikipedia.org/wiki/Independent_set_(graph_theory))

<sup>6</sup>[https://en.wikipedia.org/wiki/Travelling\\_salesman\\_problem](https://en.wikipedia.org/wiki/Travelling_salesman_problem)

*Proof.* The first inclusion,  $\mathbf{P} \subseteq \mathbf{NP}$ , is trivial if we consider that definition 7 reduces to definition 6 if we set  $q(n) = 0$ , i.e., we accept an empty certificate: languages in  $P$  are polynomially verifiable without need of certificates.

For the second inclusion, to decide  $L \in \mathbf{NP}$  we just need a machine that iterates through all certificates  $c \in \Sigma^{q(|x|)}$  and runs every time  $\mathcal{M}(x, c)$ , accepting  $x$  as soon as a pair  $(x, c)$  is accepted or rejecting  $x$  when all certificates are exhausted. This can be done in some adequately large time  $2^{n^k}$ .  $\square$

### 2.3.1 Non-deterministic Turing Machines

An alternative definition of  $\mathbf{NP}$  highlights the meaning of the class name, and will be very useful in the future.

**Definition 8.** *A non-deterministic Turing Machine (NDTM) is a TM with two different, independent transition functions. At each step, the NDTM makes an arbitrary choice as to which function to apply. Every sequence of choices defines a possible computation of the NDTM. We say that the NDTM accepts an input  $x$  if at least one computation (i.e., one of the possible arbitrary sequences of choices) terminates in an accepting state.*

There are many different ways of imagining a NDTM: one that flips a coin at each step, one that always makes the right choice towards acceptance, one that “doubles” at each step following both choices at once. Note that, while a normal, deterministic TM is a viable computational model, a NDTM is not, and has no correspondence to any current or envisionable computational device<sup>7</sup>.

Alternate definitions might refer to machines with more than two choices, with a subset of choices for every input, and so on, but they are all functionally equivalent.

We can define the class  $\mathbf{NTIME}(f)$  as the NDTM equivalent of class  $\mathbf{DTIME}(f)$ , just by replacing the TM in Definition 5 with a NDTM:

**Definition 9.** *Let  $f : \mathbb{N} \rightarrow \mathbb{N}$  be any computable function. We say that a language  $L \subseteq \Sigma^*$  is of class  $\mathbf{NTIME}(f)$ , and write  $L \in \mathbf{NTIME}(f)$ , if there is a NDTM  $\mathcal{M}$  that decides  $L$  and its worst-case time, as a function of input size, is dominated by  $f$ :*

$$L \in \mathbf{DTIME}(f) \iff \exists \mathcal{M} : L(\mathcal{M}) = L \wedge T_{\mathcal{M}} = O(f).$$

Indeed, the names “DTIME” and “NTIME” refer to the deterministic and non-deterministic reference machine. Also, the name  $\mathbf{NP}$  means “non-deterministically polynomial (time),” as the following theorem implies by setting a clear parallel between the definition of  $\mathbf{P}$  and  $\mathbf{NP}$ :

**Theorem 11.**

$$\mathbf{NP} = \bigcup_{k=0}^{\infty} \mathbf{NTIME}(n^k).$$

*Proof.* See also Theorem 2.6 in the online draft of Arora-Barak. We can prove the two inclusions separately.

Let  $L \in \mathbf{NP}$ , as in Definition 7. We can define a NDTM  $\mathcal{N}$  that, given input  $x$ , starts by non-deterministically appending a certificate  $c \in \Sigma^{q(|x|)}$ : every computation generates a different certificate. After this non-deterministic part, we run the machine  $\mathcal{M}$  from Definition 7 on the tape containing  $(x, c)$ . If  $x \in L$ , then at least one computation has written the correct certificate, and thus ends in an accepting state. On the other hand, if  $x \notin L$  then no certificate can end in acceptance. Therefore,  $\mathcal{N}$  accepts  $x$  if and only if  $x \in L$ . The NDTM  $\mathcal{N}$  performs  $q(|x|)$  steps to write the (non-deterministic) certificate, followed by the  $p(|x|)$  steps due to the execution of  $\mathcal{M}$ , and is therefore polynomial with respect to the input. Thus,  $L \in \mathbf{NTIME}(n^k)$  for some  $k \in \mathbb{N}$ .

Conversely, let  $L \in \mathbf{NTIME}(n^k)$  for some  $k \in \mathbb{N}$ . This means that  $x$  can be decided by a NDTM  $\mathcal{N}$  in time  $q(|x|) = O(|x|^k)$ , during which it performs  $q(|x|)$  arbitrary binary choices. Suppose that

<sup>7</sup>Not even quantum computing, no matter what popular science magazines write.

$x \in L$ , then there is an accepting computation by  $\mathcal{N}$ . Let  $c \in \{0, 1\}^{q(|x|)}$  be the sequence of arbitrary choices done by the accepting computation of  $\mathcal{N}(x)$ . We can use  $c$  as a certificate in Definition 7, by creating a deterministic TM  $\mathcal{M}$  that uses  $c$  to emulate  $\mathcal{N}(x)$ 's accepting computation by performing the correct choices at every step. If  $x \notin L$ , then no computation by  $\mathcal{N}(x)$  ends by accepting the input, therefore all certificates fail, and  $\mathcal{M}(x, c) = 0$  for every  $c$ . Thus, all conditions in Definition 7 hold, and  $L \in \mathbf{NP}$ .  $\square$

## 2.4 Reductions and hardness

Nobody knows if  $\mathbf{NP}$  is a proper superset of  $\mathbf{P}$ , yet. In order to better assess the problem, we need to set up a hierarchy within  $\mathbf{NP}$  in order to identify, if possible, languages that are harder than others. To do this, we resort again to *reductions*.

**Definition 10.** *Given two languages  $L, L' \in \mathbf{NP}$ , we say that  $L$  is polynomially reducible to  $L'$ , and we write  $L <_P L'$ , if there is a function  $R: \Sigma^* \rightarrow \Sigma^*$  such that*

$$x \in L \iff R(x) \in L'.$$

In other words,  $R$  maps strings in  $L$  to strings in  $L'$  and strings that are not in  $L$  to strings that are not in  $L'$ . Note that we require  $R$  to be computable in polynomial time, i.e., there must be a polynomial  $p(n)$  such that  $R(x)$  is computed in at most  $p(|x|)$  steps. If  $L <_P L'$ , we say that  $L'$  is at least as hard as  $L$ . In fact, if we have a procedure to decide  $L'$ , we can apply it to decide also  $L$  with “just” a polynomial overhead due to the reduction.

### 2.4.1 Example: Boolean formulas and the conjunctive normal form

Given  $n$  boolean variables  $x_1, \dots, x_n$ , we can define the following:

- a *term*, or *literal*, is a variable  $x_i$  or its negation  $\neg x_i$ ;
- a *clause* is a disjunction of terms;
- finally, a *formula* is a conjunction of clauses.

**Definition 11** (Conjunctive normal form). *A formula  $f$  is said to be in conjunctive normal form with  $n$  variables and  $m$  clauses if it can be written as*

$$f(x_1, \dots, x_n) = \bigwedge_{i=1}^m \bigvee_{j=1}^{l_i} g_{ij},$$

where clause  $i$  has  $l_i$  terms, every literal  $g_{ij}$  is in the form  $x_k$  or in the form  $\neg x_k$ .

For instance, the following is a CNF formula with  $n = 5$  variables and  $m = 4$  clauses:

$$\begin{aligned} f(x_1, x_2, x_3, x_4, x_5) = & (x_1 \vee \neg x_2 \vee x_4 \vee x_5) \wedge (x_2 \vee \neg x_3 \vee \neg x_4) \\ & \wedge (\neg x_1 \vee \neg x_2 \vee x_3 \vee \neg x_5) \wedge (\neg x_3 \vee \neg x_4 \vee x_5). \end{aligned} \quad (2.4)$$

Asking about the satisfiability of a CNF formula  $f$  amount at asking for a truth assignment such that every clause has at least one true literal. For example, the following assignment, among many others, satisfies (2.4):

$$x_1 = x_2 = \text{true}; \quad x_3 = x_4 = x_5 = \text{false}.$$

We can therefore say that  $f \in \text{SAT}$ .

**Definition 12** ( $k$ -CNF). *If all clauses of a CNF formula have at most  $k$  literals in them, then we say that the formula is  $k$ -CNF (conjunctive normal form with  $k$ -literal clauses).*



For instance, (2.4) is 4-CNF and, in general,  $k$ -CNF for all  $k \geq 4$ . It is not 3-CNF because it has some 4-literal clauses. Sometimes, the definition of  $k$ -CNF is stricter, and requires that every clause has *precisely*  $k$  literals. Nothing changes, since we can always write the same literal twice in order to fill the clause up.

**Definition 13.** *Given  $k \in \mathbb{N}$ , the language  $k$ -SAT is the set of all (encodings of) satisfiable  $k$ -CNF formulas.*

Let us start with a “trivial” theorem:

**Theorem 12.** *Given  $k \in \mathbb{N}$ ,*

$$k\text{-SAT} <_P \text{SAT}.$$

*Proof.* Define the reduction  $R(x)$  as follows: given a string  $x$ , if it encodes a  $k$ -CNF formula, then leave it as it is; otherwise, return an unsatisfiable formula.  $\square$

The simple reduction takes into account the fact that  $k\text{-SAT} \subseteq \text{SAT}$ , therefore if we are able to decide SAT, we can a fortiori decide  $k$ -SAT.

The following fact is less obvious:

**Theorem 13.**

$$\text{SAT} <_P 3\text{-SAT}.$$

*Proof.* Let  $f$  be a CNF formula. Suppose that  $f$  is not 3-CNF. Let clause  $i$  have  $l_i > 3$  literals:

$$\bigvee_{j=1}^{l_i} g_{ij} \tag{2.5}$$

Let us introduce a new variable,  $h$ , and split the clause as follows,

$$\left( h \vee \bigvee_{j=1}^{l_i-2} g_{ij} \right) \wedge (\neg h \vee g_{i,l_i-1} \vee g_{i,l_i}), \tag{2.6}$$

by keeping all literals, apart from the last two, in the first clause, and putting the last two in the second one. By construction, the truth assignments that satisfy (2.5) also satisfy (2.6), and viceversa. In fact, if (2.5) is satisfied then at least one of its literals are true; but then one of the two clauses of (2.6) is satisfied by the same literal, while the other can be satisfied by appropriately setting the value of the new variable  $h$ . Conversely, if both clauses in (2.6) are satisfied, then at least one of the literals in (2.5) is true, because the truth value of  $h$  alone cannot satisfy both clauses.

The step we just described transforms an  $l_i$ -literal clause into the conjunction of an  $(l_i - 1)$ -literal clause and a 3-literal clause which is satisfiable if and only if the original one was; by applying it recursively, we end up with a 3-CNF formula which is satisfiable if and only if the original  $f$  was.  $\square$

As an example, the 4-CNF formula (2.4) can be reduced to the following 3-CNF with the two additional variables  $h$  and  $k$  used to split its two 4-clauses:

$$\begin{aligned} f'(x_1, x_2, x_3, x_4, x_5, h, k) = & (h \vee x_1 \vee \neg x_2) \wedge (\neg h \vee x_4 \vee x_5) \wedge (x_2 \vee \neg x_3 \vee \neg x_4) \\ & \wedge (k \vee \neg x_1 \vee \neg x_2) \wedge (\neg k \vee x_3 \vee \neg x_5) \wedge (\neg x_3 \vee \neg x_4 \vee x_5). \end{aligned} \tag{2.7}$$

Theorem 13 is interesting because it asserts that a polynomial-time algorithm for 3-SAT would be enough for the more general problem. With the addition of Theorem 12, we can conclude that all  $k$ -SAT languages, for  $k \geq 3$ , are equivalent to each other and to the more general SAT.

On the other hand, it can be shown that  $2\text{-SAT} \in \mathbf{P}$ .




Figure 2.1: Reduction of the 3-CNF formula (2.7) to a graph for INDSET.

## 2.4.2 Example: reducing 3-SAT to INDSET

Let us see an example of reduction between two problems coming from different domains: boolean logic and graphs.

**Theorem 14.**

$$3\text{-SAT} <_P \text{INDSET}.$$

*Proof.* Let  $f$  be a 3-CNF formula. We need to transform it into a graph  $G$  and an integer  $k$  such that  $G$  has an independent set of size  $k$  if and only if  $f$  is satisfiable.

Let us represent each of the  $m$  clauses in  $f$  as a separate triangle (i.e., three connected vertices) of  $G$ , and let us label each vertex of the triangle as one of the clause's literals. Therefore,  $G$  contains  $3m$  vertices organized in  $m$  triangles.

Next, connect every vertex labeled as a variable to all vertices labeled as the corresponding negated variable: every vertex labeled “ $x_1$ ” must be connected to every vertex labeled “ $\neg x_1$ ” and so on. Fig. 2.1 shows the graph corresponding to the 3-CNF formula (2.7): each bold-edged triangle corresponds to one of the six clauses, with every node labeled with one of the literals. The dashed edges connect every literal with its negations.

It is easy to see that the original 3-CNF formula is satisfiable if and only if the graph contains an independent set of size  $k = m$  (number of clauses). Given the structure of the graph, no more than one node per triangle can appear in the independent set (nodes in the same triangle are not independent),

and if a literal appears in the independent set, then its negation does not (they would be connected by an edge, thus not independent). If the independent set has size  $m$ , then we are ensured that one literal per clause can be made true without contradictions. As an example, the six green nodes in Fig. 2.1 form an independent set and correspond to a truth assignment that satisfies  $f$ .  $\square$

## Part II

# Questions and exercises

# Appendix A

## Self-assessment questions

This chapter collects a few questions that students can try answering to assess their level of preparation.

### A.1 Computability

#### A.1.1 Recursive and recursively enumerable sets

1. Why is every finite set recursive?  
(Hint: we need to check whether  $n$  is in a finite list)
2. Try to prove that if a set is recursive, then its complement is recursive too.  
(Hint: invert 0 and 1 in the decision function's answer)
3. Let  $S$  be a recursively enumerable set, and let algorithm  $\mathcal{A}$  enumerate all elements in  $S$ . Prove that, if  $\mathcal{A}$  lists the elements of  $S$  in increasing order, then  $S$  is recursive.  
(Hint: what if  $n \notin S$ ? Is there a moment when we are sure that  $n$  will never be listed by  $\mathcal{A}$ ?)

#### A.1.2 Turing machines

1. Why do we require a TM's alphabet  $\Sigma$  and state set  $Q$  to be finite, while we accept the tape to be infinite?
2. What is the minimum size of the alphabet to have a useful TM? What about the state set?
3. Try writing machines that perform simple computations or accept simply defined strings.

# Appendix B

## Exercises

### Preliminary observations

Since the size of the alphabet, the number of tapes or the fact that they are infinite in one or both directions have no impact on the capabilities of the machine and can emulate each other, unless the exercise specifies some of these details, students are free to make their choices.

As for accepting or deciding a language, many conventions are possible. The machine may:

- erase the content of the tape and write a single “1” or “0”;
- write “1” or “0” and then stop, without bothering to clear the tape, with the convention that acceptance is encoded in the last written symbol;
- have two halting states, `halt-yes` and `halt-no`;
- any other unambiguous convention;

with the only provision that the student writes it down in the exercise solution.

### Exercise 1

For each of the following classes of Turing machines, decide whether the halting problem is computable or not. If it is, outline a procedure to compute it; if not, prove it (usually with with a reduction from the general halting problem). Unless otherwise stated, always assume that the non-blank portion of the tape is bounded, so that the input can always be finitely encoded if needed.

**1.1)** TMs with 2 symbols and at most 2 states (plus the halting state), starting from an empty (all-blank) tape.

Note: in early versions the tape wasn't mentioned. My bad.

**1.2)** TMs with at most 100 symbols and 1000000 states.

**1.3)** TMs that only move right;

**1.4)** TMs with a circular, 1000-cell tape.

**1.5)** TMs whose only tape is read-only (i.e., they always overwrite a symbol with the same one);

Suggerimento — *Actually, only one of these cases is uncomputable...*

### Solution 1

The following are minimal answers that would guarantee a good evaluation on the test.

**1.1)** The definition of the machine meet the requirements for the Busy Beaver game; Since we know the BB for up to 4 states, it means that every 2-state, 2-symbol machine has been analyzed on an empty tape, and its behavior is known. Therefore the HP is computable for this class of machines.

**1.2)** As we have seen in the lectures, 100 symbols and 1,000,000 states are much more than those needed to build a universal Turing machine  $\mathcal{U}$ . If this problem were decidable by a machine, say  $\mathcal{H}_{1,000,000}$ , then we could solve the general halting problem “does  $\mathcal{M}$  halt on input  $s$ ” by asking  $\mathcal{H}_{1,000,000}$  whether  $\mathcal{U}$  would halt on input  $(M, s)$  or not. In other words, we could reduce the general halting problem to it, therefore it is undecidable.

**1.3)** If the machine cannot visit the same cell twice, the symbol it writes won't have any effect on its future behavior. Let us simulate the machine; if it halts, then we output 1. Otherwise, sooner or later the machine will leave on its left all non-blank cells of the tape: from now on, it will only see blanks, therefore its behavior will only be determined by its state. Take into account all states entered after this moment; as soon as a state is entered for the second time, we are sure that the machine will run forever, because it is bound to repeat the same sequence of states over and over, and we can interrupt the simulation and output 0; if, on the other hand, the machine halts before repeating any state, we output 1.

**1.4)** As it has a finite alphabet and set of states (as we know from definition), the set of possible configurations of a TM with just 1000 cells is fully identified by (i) the current state, (ii) the current position, and (iii) the symbols on the tape, for a total of  $|Q| \times 1000 \times |\Sigma|^{1000}$  configurations. While this is an enormous number, a machine running indefinitely will eventually revisit the same configuration twice. So we just need to simulate a run of the machine: as soon as a configuration is revisited, we can stop simulating the machine and return 0. If, on the other hand, the simulation reaches the halt state, we can return 1.

**1.5)** Let  $n = |Q|$  be the number of states of the machine. Let us number the cells with consecutive integer numbers, and consider the cells  $a$  and  $b$  that delimit the non-null portion of the tape. Let us simulate the machine. If the machine reaches cell  $a - (n + 1)$  or  $b + n + 1$ , we will know that the machine must have entered some state twice while in the blank portion, therefore it will go on forever: we can stop the simulation and return 0. If, on the other hand, the machine always remains between cell  $a - n$  and  $b + n$ , then it will either halt (then we return 1) or revisit some already visited configuration in terms of current cell and state; in such case we know that the machine won't stop because it will deterministically repeat the same steps over and over: we can then stop the simulation and return 0.

### Exercise 2

For each of the following properties of TMs, say whether it is semantic or syntactic, and prove whether it is decidable or not.

- 2.1)  $M$  recognizes all words with an ‘a’ in them.
- 2.2)  $M$  always halts within 100 steps.
- 2.3)  $M$  either halts within 100 steps or never halts.
- 2.4)  $M$  accepts all words from the 2018 edition of the Webster’s English Dictionary.
- 2.5)  $M$  never halts in less than 100 steps.
- 2.6)  $M$  is a Turing machine.
- 2.7)  $M$  recognizes strings that encode a Turing machine (according to some predefined encoding scheme).
- 2.8)  $M$  is a TM with at most 100 states.

### Solution 2

2.1) The property is semantic, since it does not depend on the specific machine but only on the language that it recognizes. The property is also non-trivial (it can be true for some machines, false for others), therefore it satisfies the hypotheses of Rice’s theorem. We can safely conclude that it is uncomputable.

*Note: the language “All words with an ‘a’ in them” is computable. What we are talking about here is the “language” of all Turing machines that recognize it.*

2.2) Since we can always add useless states to a TM, given a machine  $M$  that satisfies the property, we can always modify it into a machine  $M'$  such that  $L(M) = L(M')$ , but that runs for more than 100 steps. Therefore the property is not semantic. It is also decidable: in order to halt within 100 steps, the machine will never visit more than 100 cells of the tape in either direction, therefore we “just” need to simulate it for at most 100 steps on all inputs of size at most 200 (a huge but finite number) and see whether it always halts within that term or not.

2.3) Again, the property is not semantic: different machines may recognize the same language but stop in a different number of steps. In this case, it is clearly undecidable: just add 100 useless states at the beginning of the execution and the property becomes “ $M$  never halts”.

2.4) The property is semantic, since it only refers to the language recognized by the machine, and is clearly non-trivial. Therefore it satisfies Rice’s Theorem hypotheses and is uncomputable. *Note: as in point 2.1, the language “all words in Webster’s” is computable, but we aren’t able to always decide whether a TM recognizes it or not.*

2.5) This is the complement of property 2.2, therefore not semantic and decidable.

2.6) The property is trivial, since all TMs trivially have it. Therefore, it is decidable by the TM that always says “yes” with no regard for the input.

2.7) The property is semantic because it refers to a specific language (strings encoding TMs). It is not trivial: even if the encoding allowed for all strings to be interpreted as a Turing machine, the only machines that possess the property would be those that recognize every string.

2.8) Deciding whether a machine has more or less than 100 states is clearly computable by just scanning the machine’s definition and counting the number of different states. The property is not semantic.



