

**COMPUTATIONAL COMPLEXITY**

Lecture notes

Ashley Montanaro, DAMTP Cambridge

am994@cam.ac.uk

**Contents**

<b>1</b>	<b>Complementary reading and acknowledgements</b>	<b>2</b>
<b>2</b>	<b>Motivation</b>	<b>3</b>
<b>3</b>	<b>The Turing machine model</b>	<b>4</b>
<b>4</b>	<b>Efficiency and time-bounded computation</b>	<b>13</b>
<b>5</b>	<b>Certificates and the class NP</b>	<b>21</b>
<b>6</b>	<b>Some more NP-complete problems</b>	<b>27</b>
<b>7</b>	<b>The P vs. NP problem and beyond</b>	<b>33</b>
<b>8</b>	<b>Space complexity</b>	<b>39</b>
<b>9</b>	<b>Randomised algorithms</b>	<b>45</b>
<b>10</b>	<b>Counting complexity and the class #P</b>	<b>53</b>
<b>11</b>	<b>Complexity classes: the summary</b>	<b>60</b>
<b>12</b>	<b>Circuit complexity</b>	<b>61</b>
<b>13</b>	<b>Decision trees</b>	<b>70</b>
<b>14</b>	<b>Communication complexity</b>	<b>77</b>
<b>15</b>	<b>Interactive proofs</b>	<b>86</b>
<b>16</b>	<b>Approximation algorithms and probabilistically checkable proofs</b>	<b>94</b>

**Health warning:** There are likely to be changes and corrections to these notes throughout the course. For updates, see <http://www.qi.damtp.cam.ac.uk/node/251/>.

Most recent changes: corrections to proof of correctness of Miller-Rabin test (Section 9.1) and minor terminology change in description of SUBSET SUM problem (Section 6).

Version 1.14 (May 28, 2013).

# 1 Complementary reading and acknowledgements

Although this course does not follow a particular textbook closely, the following books and additional resources may be particularly helpful.

- *Computational Complexity: A Modern Approach*, Sanjeev Arora and Boaz Barak. Cambridge University Press.

Encyclopaedic and recent textbook which is a useful reference for almost every topic covered in this course (a first edition, so beware typos!). Also contains much more material than we will be able to cover in this course, for those who are interested in learning more. A draft is available online at <http://www.cs.princeton.edu/theory/complexity/>.

- *Computational Complexity*, Christos Papadimitriou. Addison-Wesley.

An excellent, mathematically precise and clearly written reference for much of the more classical material in the course (especially complexity classes).

- *Introduction to the Theory of Computation*, Michael Sipser. PWS Publishing Company.

A more gentle introduction to complexity theory. Also contains a detailed introduction to automata theory, which we will not discuss in this course.

- *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Michael Garey and David Johnson. W. H. Freeman.

The standard resource for NP-completeness results.

- *Introduction to Algorithms*, Thomas Cormen, Charles Leiserson, Ronald Rivest and Clifford Stein. MIT Press and McGraw-Hill.

A standard introductory textbook to algorithms, known as CLRS.

- There are many excellent sets of lecture notes for courses on computational complexity available on the Internet, including a course in the Computer Lab here at Cambridge<sup>1</sup> and a course at the University of Maryland by Jonathan Katz<sup>2</sup>.

These notes have benefited substantially from the above references (the exposition here sometimes follows one or another closely), and also from notes from a previous course on computational complexity by Richard Jozsa.

Thanks to those who have commented on previous versions of these notes, including Raphaël Clifford, Karen Habermann, Ben Millwood and Nicholas Teh.

---

<sup>1</sup><http://www.cl.cam.ac.uk/teaching/1112/Complexity/>

<sup>2</sup><http://www.cs.umd.edu/~jkatz/complexity/f11/>

## 2 Motivation

Computational complexity theory is the study of the intrinsic difficulty of computational problems. It is a young field by mathematical standards; while the foundations of the theory were laid by Alan Turing and others in the 1930's, computational complexity in its modern sense has only been an object of serious study since the 1970's, and many of the field's most basic and accessible questions remain open. The sort of questions which the theory addresses range from the immediately practical to the deeply philosophical, including:

- Are there problems which cannot be solved by computer?
- Is there an efficient algorithm for register allocation in a computer's central processing unit?
- Can we automate the process of discovering mathematical theorems?
- Can all algorithms be written to run efficiently on a computer with many processors?
- How quickly can we multiply two  $n \times n$  matrices?
- Are some types of computational resource more powerful than others?

In general, we will be concerned with the question of which problems we can solve, given that we have only limited resources (e.g. restrictions on the amount of time or space available to us). Rather than consider individual problems (e.g. "is the Riemann hypothesis true?"), it will turn out to be more fruitful to study families of problems of growing size (e.g. "given a mathematical claim which is  $n$  characters long, does it have a proof which is at most  $n^2$  characters long?"), and study how the resources required to solve the problem in question grow with its size. This perspective leads to a remarkably intricate and precise classification of problems by difficulty.

### 2.1 Notation

We collect here some notation which will be used throughout the course.  $\Sigma$  will usually denote an *alphabet*, i.e. a finite set of symbols. Notable alphabets include the binary alphabet  $\{0, 1\}$ ; an element  $x \in \{0, 1\}$  is known as a bit. A string over  $\Sigma$  is an ordered sequence (possibly empty) of elements of  $\Sigma$ . For strings  $\sigma$ ,  $|\sigma|$  denotes the length of  $\sigma$ .  $\Sigma^k$  denotes the set of length  $k$  strings over  $\Sigma$ , and  $\Sigma^*$  denotes the set of all finite length strings over  $\Sigma$ , i.e.  $\Sigma^* = \bigcup_{k \geq 0} \Sigma^k$ . For  $\sigma \in \Sigma$ , the notation  $\sigma^k$  denotes the string consisting of  $k$   $\sigma$ 's (e.g.  $0^3 = 000$ ).

We will sometimes need to write strings in lexicographic order. This is the ordering where all strings are first ordered by length (shortest first), then strings of length  $n$  are ordered such that  $a \in \Sigma^n$  comes before  $b \in \Sigma^n$  if, for the first index  $i$  on which  $a_i \neq b_i$ ,  $a_i$  comes before  $b_i$  in  $\Sigma$ . For example, lexicographic ordering of the set of nonempty binary strings is  $\{0, 1, 00, 01, 10, 11, 000, \dots\}$ .

We use the notation  $A \subseteq B$  (resp.  $A \subset B$ ) to imply that  $A$  is contained (resp. strictly contained) within  $B$ .  $A \Delta B$  is used for the symmetric difference of sets  $A$  and  $B$ , i.e.  $(A \cup B) \setminus (A \cap B)$ . We let  $[n]$  denote the set  $\{1, \dots, n\}$ . For  $x \geq 0$ ,  $\lfloor x \rfloor$  and  $\lceil x \rceil$  denote the floor and ceiling of  $x$  respectively, i.e.  $\lfloor x \rfloor = \max\{n \in \mathbb{Z} : n \leq x\}$ ,  $\lceil x \rceil = \min\{n \in \mathbb{Z} : n \geq x\}$ .

We will often need to encode elements  $s \in S$ , for some set  $S$ , as binary strings. We write  $\underline{s}$  for such an encoding of  $s$ , which will usually be done in some straightforward manner. For example, integers  $x \in \mathbb{N}$  are represented as  $\underline{x} \in \{0, 1\}^*$  by just writing the digits of  $x$  in binary.



Figure 1: A typical configuration of a Turing machine, and the starting configuration on input 21.

### 3 The Turing machine model

How are we to model the idea of computation? For us a computation will be the execution of an *algorithm*. An algorithm is a method for solving a problem; more precisely, an effective method for calculating a function, expressed as a finite list of well-defined instructions. We will implement our algorithms using a basic, yet surprisingly powerful, kind of computer called a Turing machine.

A Turing machine can be thought of as a physical computational device consisting of a *tape* and a *head*. The head moves along the tape, scanning and modifying its contents according to the currently scanned symbol and its internal state. A “program” for a Turing machine specifies how it should operate at each step of its execution (i.e. in which direction it should move and how it should update the tape and its internal state). We assume that the tape is infinite in one direction and the head starts at the leftmost end.

Formally, a Turing machine  $M$  is specified by a triple  $(\Sigma, K, \delta)$ , where:

- $\Sigma$  is a finite set of symbols, called the *alphabet* of  $M$ . We assume that  $\Sigma$  contains two special symbols  $\square$  and  $\triangleright$ , called the *blank* and *start* symbols, respectively.
- $K$  is a finite set of *states*. We assume that  $K$  contains a designated start state **START**, and a designated halting state **HALT**.  $K$  describes the set of possible “mental states” of  $M$ .
- $\delta$  is a function such that  $\delta : K \times \Sigma \rightarrow K \times \Sigma \times \{\leftarrow, -, \rightarrow\}$ .  $\delta$  describes the behaviour of  $M$  and is called the transition function of  $M$ .

A tape is an infinite sequence of cells, each of which holds an element of  $\Sigma$ . The tape initially contains  $\triangleright$ , followed by a finite string  $x \in (\Sigma \setminus \{\square\})^k$  called the *input*, followed by an infinite number of blank cells  $\square\square\dots$ . For brevity, in future we will often suppress the trailing infinite sequence of  $\square$ 's when we specify what the tape contains.

At any given time, the current state of the computation being performed by  $M$  is completely specified by a *configuration*. A configuration is given by a triple  $(\ell, q, r)$ , where  $q \in K$  is a state and specifies the current “state of mind” of  $M$ , and  $\ell, r \in \Sigma^*$  are strings, where  $\ell$  describes the tape to the left of the head (including the current symbol scanned by the head, which is thus the last element in  $\ell$ ) and  $r$  describes the tape to the right of the head. Thus the initial configuration of  $M$  on input  $x$  is  $(\triangleright, \text{START}, x)$ .

At each step of  $M$ 's operation, it updates its configuration according to  $\delta$ . Assume that at a given step  $M$  is in state  $q$ , and the symbol currently being scanned by the head is  $\sigma$ . Then, if  $\delta(q, \sigma) = (q', \sigma', d)$ , where  $d \in \{\leftarrow, -, \rightarrow\}$ :

- the symbol on the tape at the position currently scanned by the head is replaced with  $\sigma'$ ;
- the state of  $M$  is replaced with  $q'$ ;

- if  $d = \leftarrow$ , the head moves one place to the left; if  $d = -$ , the head stays where it is; if  $d = \rightarrow$ , the head moves one place to the right.

We assume that, for all states  $p$  and  $q$ , if  $\delta(p, \triangleright) = (q, \sigma, d)$  then  $\sigma = \triangleright$  and  $d \neq \leftarrow$ . That is, the head never falls off the left end of the tape and never erases the start symbol. We also assume that  $\delta(\text{HALT}, \sigma) = (\text{HALT}, \sigma, -)$ . That is, when  $M$  is in the halt state, it can no longer modify its configuration. If  $M$  has halted, we call the contents of the tape the output (excluding the  $\triangleright$  symbol on the left of the tape and the infinite string of  $\square$ 's on the right). If  $M$  halts with output  $y$  on input  $x$ , we write  $M(x) = y$ , and write  $M : \Sigma^* \rightarrow \Sigma^*$  for the function computed by  $M$ . Of course, some machines  $M$  may not halt at all on certain inputs  $x$ , but just run forever. If  $M$  does not halt on input  $x$ , we can write  $M(x) = \nearrow$  (in fact, we will never need to do so).

### 3.1 Example

We describe a Turing machine  $M$  below which computes the function  $\text{RMZ} : \{0, 1\}^* \rightarrow \{0, 1\}^*$ , where  $\text{RMZ}(x)$  is equal to  $x$  with the rightmost 0 changed to a 1. If there are no 0's in  $x$ , the machine outputs  $x$  unchanged.  $M$  operates on alphabet  $\{\triangleright, \square, 0, 1\}$ , has states  $\{\text{START}, \text{TOEND}, \text{FINDZERO}, \text{HALT}\}$ , and has transition function  $\delta$  which performs the following maps.

$$\begin{aligned}
(\text{START}, \triangleright) &\mapsto (\text{TOEND}, \triangleright, \rightarrow) \\
(\text{TOEND}, 0) &\mapsto (\text{TOEND}, 0, \rightarrow) \\
(\text{TOEND}, 1) &\mapsto (\text{TOEND}, 1, \rightarrow) \\
(\text{TOEND}, \square) &\mapsto (\text{FINDZERO}, \square, \leftarrow) \\
(\text{FINDZERO}, 0) &\mapsto (\text{HALT}, 1, -) \\
(\text{FINDZERO}, 1) &\mapsto (\text{FINDZERO}, 1, \leftarrow) \\
(\text{FINDZERO}, \triangleright) &\mapsto (\text{HALT}, \triangleright, \rightarrow)
\end{aligned}$$

For brevity, transitions that can never occur are not specified above; however, it should already be clear that writing programs in the Turing machine model directly is a tedious process. In future, we will usually describe algorithms in this model more informally.

### 3.2 Observations about Turing machines

- **Alphabet.** We immediately observe that we can fix the alphabet  $\Sigma = \{\triangleright, \square, 0, 1\}$  (i.e. use a binary alphabet, with two additional special symbols) without restricting the model, because we can encode any more complicated alphabets in binary. Imagine that we have a machine  $M$  whose alphabet  $\Sigma$  contains  $K$  symbols (other than  $\triangleright, \square$ ). Encode the  $i$ 'th symbol as a  $k := \lceil \log_2 K \rceil$ -bit binary number, i.e. a string of  $k$  0's and 1's. Define a new Turing machine  $\widetilde{M}$  which will simulate  $M$  as follows.

The tape on which  $\widetilde{M}$  operates contains the binary encodings of the symbols of the tape on which  $M$  operates, in the same order. In order to simulate one step of  $M$ , the machine  $\widetilde{M}$ :

1. Reads the  $k$  bits from its tape which encode the current symbol of  $\Sigma$  which  $M$  is scanning, using its state to store what this current symbol is. If the bits read do not correspond to a symbol of  $\Sigma$ ,  $\widetilde{M}$  can just halt.
2. Uses  $M$ 's transition function to decide what the next state of  $M$  should be, given this input symbol, and stores this information in its own state.

- Updates the  $k$  bits on the tape to encode the new symbol which  $M$  would write to the tape.

This gives a very simple example of an encoding, i.e. a map which allows us to translate between different alphabets. Henceforth, we will often assume that our Turing machines operate on a binary alphabet. Similarly, we can assume that the states of the Turing machine are given by the integers  $1, \dots, m$ , for some finite  $m$ .

- **Encoding of the input.** We may sometimes wish to pass multiple inputs  $x_1, \dots, x_k$  to  $M$ . To do so, we can simply introduce a new symbol “,” to the alphabet, enabling  $M$  to determine when one input ends and another begins.
- **Representation and countability.** A Turing machine  $M$ 's behaviour is completely specified by its transition function and the identity of the special states and symbols  $\square, \triangleright, \text{START}, \text{HALT}$ . Thus  $M$  can be represented by a finite sequence of integers, or (equivalently) a finite length bit-string. There are many reasonable ways that this representation can be implemented, and the details are not too important. Henceforth we simply assume that we have fixed one such representation method, and use the notation  $\underline{M}$  for the bit-string representing  $M$ . The existence of such a representation implies that the set of Turing machines is countable, i.e. in bijection with  $\mathbb{N}$ .

It will be convenient later on to assume that every string  $x \in \{0, 1\}^*$  actually corresponds to some Turing machine. This can be done, for example, by associating each string  $\underline{M}$  which is not a valid encoding of some Turing machine  $M$  with a single fixed machine, such as a machine which immediately halts, whatever the input.

### 3.3 Big-O notation

Later on, we will need to compare running times of different Turing machines. In order to do so while eliding irrelevant details, an important tool will be asymptotic (big-O) notation, which is defined as follows. Let  $f, g : \mathbb{N} \rightarrow \mathbb{R}$  be functions. We say that:

- $f(n) = O(g(n))$  if there exist  $c > 0$  and integer  $n_0$  such that for all  $n \geq n_0$ ,  $f(n) \leq cg(n)$ .
- $f(n) = \Omega(g(n))$  if there exist  $c > 0$  and integer  $n_0$  such that for all  $n \geq n_0$ ,  $f(n) \geq cg(n)$ .
- $f(n) = \Theta(g(n))$  if there exist  $c_1, c_2 > 0$  and integer  $n_0$  such that for all  $n \geq n_0$ ,  $c_1 g(n) \leq f(n) \leq c_2 g(n)$ .
- $f(n) = o(g(n))$  if, for all  $\epsilon > 0$ , there exists integer  $n_0$  such that for all  $n \geq n_0$ ,  $f(n) \leq \epsilon g(n)$ .
- $f(n) = \omega(g(n))$  if, for all  $\epsilon > 0$ , there exists integer  $n_0$  such that for all  $n \geq n_0$ ,  $g(n) \leq \epsilon f(n)$ .

Thus:  $f(n) = O(g(n))$  if and only if  $g(n) = \Omega(f(n))$ ;  $f(n) = \Theta(g(n))$  if and only if  $f(n) = O(g(n))$  and  $f(n) = \Omega(g(n))$ ; and  $f(n) = o(g(n))$  if and only if  $g(n) = \omega(f(n))$ . The notations  $O, \Omega, \Theta, o, \omega$  can be viewed as asymptotic versions of  $\leq, \geq, =, <, >$  respectively. Beware that elsewhere in mathematics the notation  $f(n) = O(g(n))$  is sometimes used for what we write as  $f(n) = \Theta(g(n))$ . For example, in our notation  $n^2 = O(n^3)$ .

Some simple examples:




Figure 2: A typical configuration of a Turing machine with three tapes (input, work and output) on input 1012.

- If  $f(n) = 100n^3 + 3n^2 + 7$ , then  $f(n) = O(n^3)$ . In general, if  $f(n)$  is a degree  $k$  polynomial, we have  $f(n) = \Theta(n^k)$ ,  $f(n) = \omega(n^{k-1})$ ,  $f(n) = o(n^{k+1})$ . In future, we write “ $f(n) = \text{poly}(n)$ ” as shorthand for “ $f(n) = O(n^k)$  for some  $k$ ”.
- Take  $f(n) = n^k$  for some fixed  $k$ ,  $g(n) = c^n$  for some fixed  $c > 1$ . Then  $f(n) = o(g(n))$ . For example,  $n^{1000} = o(1.0001^n)$ .

We sometimes use big-O notation as part of a more complicated mathematical expression. For example,  $f(n) = 2^{O(n)}$  is shorthand for “there exist  $c > 0$  and integer  $n_0$  such that for all  $n \geq n_0$ ,  $f(n) \leq 2^{cn}$ ”. One can even write (valid!) expressions like “ $O(n) + O(n) = O(n)$ ”.

### 3.4 Multiple-tape Turing machines

The Turing machine model is very simple and there are many ways in which it can (apparently!) be generalised. Remarkably, these “generalisations” usually turn out not to be any more powerful than the original model.

A natural example of such a generalisation is to give the machine access to multiple tapes. A  $k$ -tape Turing machine  $M$  is a machine equipped with  $k$  tapes and  $k$  heads. The input is provided on a designated input tape, and the output is written on a designated (separate) output tape. The input tape is usually considered to be read-only, i.e.  $M$  does not modify it during its operation. The remaining  $k - 2$  tapes are work tapes that can be written to and read from throughout  $M$ 's operation. The work and output tapes are initially empty, apart from the start symbol. At each stage of the computation,  $M$  scans the tapes under each of its heads, and performs an action on each (modifying the tape under the head and moving to the left or right, or staying still).  $M$ 's transition function is thus of the form  $\delta : K \times \Sigma^k \rightarrow K \times (\Sigma \times \{\leftarrow, -, \rightarrow\})^k$ . Observe that  $M$ 's internal state is shared across tapes.

**Theorem 3.1.** *Given a description of any  $k$ -tape Turing machine  $M$  operating within  $T(n)$  steps on inputs of length  $n$ , we can give a single tape Turing machine  $M'$  operating within  $O(T(n)^2)$  steps such that  $M'(x) = M(x)$  for all inputs  $x$ .*

*Proof.* Let the input  $x$  be of length  $n$ . The basic idea is that our machine  $M'$  will encode the  $k$  tapes of  $M$  within one tape by storing the  $j$ 'th cell of the  $i$ 'th tape of  $M$  at position  $n + (j - 1)k + i$ . Thus the first tape is stored at  $n + 1, n + k + 1, n + 2k + 1, \dots$ , etc. The alphabet of  $M'$  will be twice the size of the alphabet of  $M$ , containing two elements  $a, \hat{a}$  for each element  $a$  of the alphabet of  $M$ . A “hat” implies that there is a head at that position. Exactly one cell in the encoding of each tape will include an element with a hat.

To start with,  $M'$  copies its input into the correct positions to encode the input tape and initialises the first block  $n + 1, \dots, n + k$  to  $\hat{\diamond}$ . This uses  $O(n^2)$  steps (note that  $k$  is constant). Then,

to simulate each step of  $M$ 's computation,  $M'$  scans the tape from left to right to determine the symbols scanned by each of the  $k$  heads, which it stores in its internal state. Once  $M'$  knows these symbols, it can compute  $M$ 's transition function and update the encodings of the head positions and tapes accordingly. As  $M$  operates within  $T(n)$  steps, it never reaches more than  $T(n)$  locations in each of its tapes, so simulating each step of  $M$ 's computation takes at most  $O(T(n))$  steps. When  $M$  halts,  $M'$  copies the contents of the output tape to the start of its tape and halts. As there are at most  $T(n)$  steps of  $M$  to simulate, the overall simulation is within  $O(T(n)^2)$  steps.  $\square$

We say that a model of computation is equivalent to the Turing machine model if it can simulate a Turing machine, and can also be simulated by a Turing machine. Thus the above theorem states that multiple-tape Turing machines are equivalent to single-tape Turing machines. The simple fact used in the proof that the amount of space used by a computation cannot be greater than the amount of time used will come up again later.

### 3.5 The universal Turing machine

The fact that Turing machines  $M$  can be represented as bit-strings  $\underline{M}$  allows  $M$  to be given as input to another Turing machine  $U$ , raising the possibility of *universal* Turing machines: Turing machines which simulate the behaviour of any possible  $M$ . We now briefly describe (without going into the technical details) how such a simulation can be done.

**Theorem 3.2.** *There exists a two-tape Turing machine  $U$  such that  $U(\underline{M}, x) = M(x)$  for all Turing machines  $M$ .*

*Proof.* As discussed in Section 3.4, we may assume that  $M$  has one tape.  $U$ 's input tape stores a description of  $M$  (which is indeed fully specified by its transition function  $\delta$  and identities of the special states START, HALT) and the input  $x$ .  $U$ 's second tape stores the current configuration  $(\ell, q, r)$  of  $M$  simply by concatenating the three elements of the configuration, separated by “,” symbols. To simulate a step of  $M$ ,  $U$  scans the input tape to find the new state, new symbol to be written and head movement to be performed.  $U$  then updates the second tape accordingly. If  $M$  halts,  $U$  cleans up the second tape to  $(\ell, r)$  and then halts.  $\square$

The overloading of the notation  $M$  highlights the triple role of Turing machines: they are simultaneously machines, functions, and data!

### 3.6 Computability

It is not the case that all functions  $f : \{0, 1\}^* \rightarrow \{0, 1\}$  can be computed by some Turing machine. Indeed, this follows from quite general arguments: as discussed, the set of Turing machines is in bijection with the countable set  $\{0, 1\}^*$ , while the set of all functions  $f : \{0, 1\}^* \rightarrow \{0, 1\}$  is in bijection with the uncountable set of all subsets of  $\{0, 1\}^*$ . Explicitly, we have the following theorem.

**Theorem 3.3.** *There exists a function  $UC : \{0, 1\}^* \rightarrow \{0, 1\}$  which is not computable by a Turing machine.*

*Proof.* Define UC as follows. For every  $\underline{M} \in \{0, 1\}^*$ , if  $M(\underline{M}) = 1$ , then  $UC(\underline{M}) = 0$ ; otherwise (i.e. if  $M(\underline{M}) \neq 1$  or  $M$  does not halt on input  $\underline{M}$ ),  $UC(\underline{M}) = 1$ . Assume towards a contradiction



	0	1	00	01	...	<u>N</u>
0	<b>0</b>	1	0	0		
1	0	<b>1</b>	0	1		
00	1	1	<b>1</b>	0		
01	0	1	0	<b>0</b>		
⋮					⋱	
<u>M</u>						1 - M( <u>N</u> )

Figure 3: Diagonalisation for proving undecidability.

that UC is computable, so there exists a Turing machine  $N$  such that, for all  $\underline{M} \in \{0, 1\}^*$ ,  $N(\underline{M}) = \text{UC}(\underline{M})$ . In particular,  $N(\underline{N}) = \text{UC}(\underline{N})$ . But we have defined

$$\text{UC}(\underline{N}) = 1 \Leftrightarrow N(\underline{N}) \neq 1,$$

so this is impossible. □

This sort of argument is known as diagonalisation, for the following reason. We write down an (infinite!) matrix  $A$  whose rows and columns are indexed by bit-strings  $x, y \in \{0, 1\}^*$ , in lexicographic order. The rows are intended to correspond to (encodings of) Turing machines, and the columns correspond to inputs to Turing machines. Define  $A_{xy} = 1$  if machine  $x$  halts with output 1 on input  $y$ , and  $A_{xy} = 0$  otherwise. Then the function  $\text{UC}(x)$  is defined by negating the diagonal of this table. Since the rows represent all Turing machines, and for all  $x$ ,  $\text{UC}(x)$  differs on the  $i$ 'th input from the function computed by the  $i$ 'th Turing machine, the function  $\text{UC}(x)$  cannot be computed for all  $x$  by a Turing machine. See Figure 3 for an illustration.

The function UC may appear fairly contrived. However, it turns out that some very natural functions are also not computable by Turing machines; the canonical example of this phenomenon is the so-called halting problem. The function  $\text{HALT}(\underline{M}, x)$  is defined as:

$$\text{HALT}(\underline{M}, x) = \begin{cases} 1 & \text{if } M \text{ halts on input } x \\ 0 & \text{otherwise.} \end{cases}$$

**Theorem 3.4.** *HALT is not computable by a Turing machine.*

*Proof.* Suppose for a contradiction that there exists a Turing machine  $M$  which computes HALT; we will show that this implies the existence of a machine  $M'$  which computes  $\text{UC}(\underline{N})$  for any  $N$ , contradicting Theorem 3.3. On input  $\underline{N}$ ,  $M'$  computes  $\text{HALT}(\underline{N}, \underline{N})$ . If the answer is 0 (i.e.  $N$  does not halt on input  $\underline{N}$ ),  $M'$  outputs 1. Otherwise,  $M'$  simulates  $N$  on input  $\underline{N}$  using Theorem 3.2, and outputs 0 if  $N$ 's output would be 1, or 1 if  $N$ 's output would not be 1. Note that this can be done in finite time because we know that  $N$  halts. □

This is our first example of a fundamental technique in computational complexity theory: proving hardness of some problem A by reducing some problem B, which is known to be hard, to solving problem A. We state without proof some other problems corresponding to functions which are now known to be uncomputable.

- Hilbert's tenth problem: given the description of a multivariate polynomial with integer coefficients, does it have an integer root?

- The Post correspondence problem. We are given a collection  $S$  of dominos, each containing two strings from some alphabet  $\Sigma$  (one on the top half of the domino, one on the bottom). For example,

$$S = \left\{ \begin{bmatrix} a \\ ab \end{bmatrix}, \begin{bmatrix} b \\ a \end{bmatrix}, \begin{bmatrix} abc \\ c \end{bmatrix} \right\}.$$

The problem is to determine whether, by lining up dominos from  $S$  (with repetitions allowed) one can make the concatenated strings on the top of the dominos equal to the concatenated strings on the bottom. For example,

$$\begin{bmatrix} a \\ ab \end{bmatrix} \begin{bmatrix} b \\ a \end{bmatrix} \begin{bmatrix} a \\ ab \end{bmatrix} \begin{bmatrix} abc \\ c \end{bmatrix}$$

would be a valid solution.

- A *Wang tile* is a unit square with coloured edges. Given a set  $S$  of Wang tiles, determine whether tiles picked from  $S$  (without rotations or reflections) can be arranged edge-to-edge to tile the plane, such that abutting edges of adjacent tiles have the same colour. For example, the following set  $S$  does satisfy this property.

$$S = \left\{ \begin{array}{c} \text{Green/Blue} \\ \text{Red/Blue} \\ \text{Green/Blue} \\ \text{Red/Blue} \end{array}, \begin{array}{c} \text{Red/Blue} \\ \text{Green/Blue} \\ \text{Red/Blue} \\ \text{Green/Blue} \end{array}, \begin{array}{c} \text{Green/Blue} \\ \text{Red/Blue} \\ \text{Green/Blue} \\ \text{Red/Blue} \end{array}, \begin{array}{c} \text{Red/Blue} \\ \text{Green/Blue} \\ \text{Red/Blue} \\ \text{Green/Blue} \end{array} \right\}.$$

Could there be other “reasonable” models of computation beyond the Turing machine which can do things that the Turing machine cannot, such as solving the halting problem? Here “reasonable” should be taken to mean: “corresponding to computations we can perform in our physical universe”. The (unprovable!) **Church-Turing thesis** is that this is not the case: informally, “everything that can be computed can be computed by a Turing machine”. Some evidence for this thesis is provided by the fact that many apparent generalisations of the Turing machine model turn out to be equivalent to Turing machines. Here we will assume the truth of the Church-Turing thesis throughout, and hence simply use the term “computable” as shorthand for “computable by a Turing machine”.

Those of you taking the Quantum Computation course next term will learn about the potential challenge to the Turing machine model posed by quantum computers, which are machines designed to use quantum mechanics in an essential manner to do things which computers based only on the laws of classical physics cannot. It turns out that quantum computers can be simulated by the (purely classical!) Turing machine. However, we do not know how to perform this simulation *efficiently*, in a certain sense. This notion of efficiency will be the topic of the next section.

### 3.7 Decision problems and languages

We will often be concerned with decision problems, i.e. problems with a yes/no answer. Such problems can be expressed naturally in terms of *languages*. A language is a subset of strings,  $\mathcal{L} \subseteq \Sigma^*$ . We say that  $\mathcal{L}$  is *trivial* if either  $\mathcal{L} = \emptyset$ , or  $\mathcal{L} = \Sigma^*$ . Let  $M$  be a Turing machine. For each input  $x$ , if  $M(x) = 1$ , we say that  $M$  accepts  $x$ ; if  $M$  halts on input  $x$  and  $M(x) \neq 1$ , we say that  $M$  rejects  $x$ . We say that  $M$  *decides*  $\mathcal{L}$  if it computes the function  $f_{\mathcal{L}} : \Sigma^* \rightarrow \{0, 1\}$ , where  $f_{\mathcal{L}}(x) = 1$  if  $x \in \mathcal{L}$ , and  $f_{\mathcal{L}}(x) = 0$  if  $x \notin \mathcal{L}$ . If there exists such an  $M$  then we say that  $\mathcal{L}$  is decidable. On the other hand, we say that  $\mathcal{L}$  is *undecidable* if  $f_{\mathcal{L}}$  is uncomputable. Theorem 3.4 above therefore says that the language

$$\text{HALT} = \{(\underline{M}, x) : M \text{ is a Turing machine that halts on input } x\}$$

is undecidable. We also say that  $M$  recognises  $\mathcal{L}$  if:

- $M(x) = 1$  for all  $x \in \mathcal{L}$ ;
- for all  $x \notin \mathcal{L}$ , either  $M$  halts with output  $\neq 1$ , or  $M$  does not halt.

Thus all decidable languages are recognisable, but the converse is not necessarily true.

### 3.8 The Entscheidungsproblem

We now informally discuss a way in which the concrete-seeming Turing machine model can be used to attack problems in the foundations of mathematics itself. The forbidding-sounding “Entscheidungsproblem” (which is simply German for “decision problem”) is the following question, first posed by Hilbert in 1928. Does there exist an algorithm which, given a set of axioms and a mathematical proposition, decides whether it is provable from the axioms? In other words, is the language of valid mathematical statements in a given axiomatic system decidable?

For example, consider statements about the natural numbers. We would like an algorithm which determines whether statements like the following are true:

$$\forall a, b, c, n[(a, b, c > 0 \wedge n > 2) \Rightarrow a^n + b^n \neq c^n].$$

We can define the alphabet of statements about the natural numbers as

$$\Sigma_{\mathbb{N}} := \{\wedge, \vee, \neg, (, ), \forall, \exists, =, <, >, +, \times, 0, 1, x\},$$

where  $x$  denotes the possibility to have variables in our statements. In order to determine whether such statements are provable, we also need to choose a set of axioms. A standard set of axioms for the natural numbers is called Peano arithmetic; the details of these are not so important for the high-level discussion here.

**Theorem 3.5.** *The language of statements about the natural numbers provable from the axioms of Peano arithmetic is undecidable.*

*Proof idea.* Let  $M$  be a Turing machine and  $w$  be a bit-string. The idea is to construct a formula  $\phi_{M,w}$  in the language of statements about the natural numbers that contains one free variable  $x$ , and such that  $\exists x \phi_{M,w}$  is true if and only if  $M$  accepts  $w$ .  $x$  is intended to encode a computation history (i.e. a complete description of the operation of a Turing machine) as an integer, and the formula  $\phi_{M,w}$  is designed to check whether  $x$  is a valid computation history for  $M$  on input  $w$ , corresponding to  $M$  accepting  $w$ . This checking can be performed using arithmetic operations  $+$ ,  $\times$ . The details of this process are quite technical, but it should at least be plausible that such an encoding can be carried out; any possible configuration of  $M$  can be encoded as an integer, and given two configurations  $c_1, c_2$ , the constraint that  $M$  maps  $c_1$  to  $c_2$  can be enforced by “local” arithmetical checks, corresponding to the locality of Turing machines.  $\square$

### 3.9 Historical notes and further reading

The Turing machine was invented by Alan Turing in 1936, while a Fellow of King’s College, Cambridge. His seminal paper “On Computable Numbers, with an Application to the Entscheidungsproblem” can easily be found online and is worth reading. Turing was actually not the first

to prove that the Entscheidungsproblem is unsolvable; Alonzo Church beat him by a few months, using a different model of computation (the so-called  $\lambda$ -calculus). In an appendix to Turing's paper, he proves that Church's model is in fact equivalent to the Turing machine, giving the first evidence for what is now called the Church-Turing thesis.

Every introductory textbook on computational complexity has its own description of the Turing machine model (each one usually explained slightly differently), e.g. Arora-Barak chapter 1, Papadimitriou chapter 2. The discussion here about the Entscheidungsproblem is essentially taken from Sipser chapter 6 (see also Arora-Barak, end of chapter 1).

Hilbert's tenth problem was proposed in 1900 but only proven undecidable in 1970. An interesting survey of undecidability in various areas of mathematics has recently been produced by Poonen<sup>1</sup>. For more on the connections between undecidability and the foundations of mathematics, including Gödel's incompleteness theorem, see Papadimitriou chapter 6 or Sipser chapter 6.

---

<sup>1</sup><http://arxiv.org/pdf/1204.0299v1.pdf>

## 4 Efficiency and time-bounded computation

It is intuitively clear that some algorithms are more efficient than others; the Turing machine model allows us to formalise this notion. The first important resource which we will consider is time. Let  $f : \Sigma^* \rightarrow \Sigma^*$ ,  $T : \mathbb{N} \rightarrow \mathbb{N}$  be functions. If  $M$  is a Turing machine with alphabet  $\Sigma$ , we say that  $M$  computes  $f$  in time  $T(n)$  if, for all  $x \in \Sigma^*$ ,  $M$  halts with output  $f(x)$  using at most  $T(|x|)$  steps. We stress that the “ $n$ ” in  $T(n)$  is used as a placeholder; there is no implication that  $|x| = n$ . In particular, we say  $M$  computes  $f$  in time  $\text{poly}(n)$  if, for all  $x \in \Sigma^*$ ,  $M$  halts with output  $f(x)$  in time  $\text{poly}(|x|)$ .

For example, the Turing machine in Section 3.1 computes the function RMZ in time  $2(n + 1)$ . Observe that this is a “worst-case” notion: for some inputs, this machine halts more quickly, but we are interested in its behaviour on the worst possible input. We first show that there is little point in calculating running times exactly for Turing machines, as one can essentially always tweak the machine to make it run faster.

**Theorem 4.1** (Linear Speedup Theorem). *For any function  $f : \Sigma^* \rightarrow \Sigma^*$ , if there is a  $k$ -tape Turing machine  $M$  which computes  $f$  in time  $T(n) \geq n$ , then for any  $\epsilon > 0$  there is a  $k'$ -tape Turing machine  $N$  which computes  $f$  in time  $\epsilon T(n) + n + 2$ . If  $k = 1$ , then  $k' = 2$ ; otherwise  $k' = k$ .*

*Proof.* For simplicity, assume in the proof that  $k = 1$ ,  $k' = 2$  (the general case is similar). We define a new Turing machine  $N$  with two tapes.  $N$ 's alphabet contains, as well as every symbol in  $M$ 's alphabet, a new symbol for each possible  $m$ -tuple of symbols in the alphabet of  $M$ , for some  $m$ . That is, if  $M$  had alphabet  $\Sigma$ ,  $N$  has alphabet  $\Sigma \cup \Sigma^m$ . The idea will be to simulate  $m$  steps of  $M$ 's operation using only one step of  $N$ . To start with,  $N$  reads its input tape. Whenever it reads a  $m$ -tuple  $(\sigma_1, \dots, \sigma_m)$ , it writes the corresponding symbol in its extended alphabet to its work tape; if a  $\square$  is encountered partway through a  $m$ -tuple (meaning  $N$  has got to the end of the input), the symbol is padded by the right number of  $\square$ 's. When  $N$  has finished reading the input  $x$  (which uses  $|x| + 2$  steps), it returns the head of its work tape to the start (using a further  $\lceil |x|/m \rceil$  steps). The work tape of  $N$  is henceforth treated as its input tape.

Now  $m$  steps of  $M$  can be simulated using 6 steps of  $N$ , as follows. First,  $N$  reads the cells under its head and the cells immediately to the left and right (by moving its head one step to the left, then two to the right, then one to the left). Dependent on the contents of these cells,  $N$  updates them according to  $m$  steps of  $M$ 's transition function, using at most two more steps. This can be done because  $m$  steps of  $M$  cannot travel further than the current cell of  $N$ , or the cells immediately to the left and right; the updates made by  $m$  steps of  $M$  can only modify the current cell of  $N$ , or one of its two neighbours, so require at most two more steps to simulate.

The overall number of steps used to simulate  $M$ 's computation on input  $x$  is thus at most  $|x| + 2 + \lceil |x|/m \rceil + 6\lceil T(|x|)/m \rceil$ . Recalling that  $T(|x|) \geq |x|$  and taking  $m = \lceil 7/\epsilon \rceil$  implies the claimed result.  $\square$

The restriction that  $T(n) \geq n$  in the Linear Speedup Theorem is not very significant, as Turing machines running in time less than this bound are generally considered uninteresting because they cannot read all of their input. The proof of the theorem crucially used the fact that computation is *local* in the Turing machine model, i.e. the head can only affect the tape around its current position. This idea will occur again later.




Figure 4: A directed graph and its representation as an adjacency matrix and an adjacency list.

## 4.1 Problems and complexity classes

A *complexity class* is a family of languages<sup>1</sup>. We write the names of complexity classes in **SANS SERIF**, and usually write the names of languages in **SMALL CAPS**. A major goal of computational complexity theory is to classify decision problems (i.e. languages; we will use the terms “decision problem” and “language” essentially interchangeably throughout) into complexity classes of similar levels of difficulty. Here are some examples of the sort of decision problems which we will consider.

- **PRIMES**. Given an integer  $N$  specified as  $n$  binary digits, is  $N$  prime? Equivalently, decide the language  $\text{PRIMES} = \{N \in \{0,1\}^* : \forall p, q \geq 2, N \neq p \times q\}$ , where  $\times$  is usual integer multiplication. Observe that we could have specified  $N$  in decimal without significantly changing the complexity of this problem.
- We will frequently be interested in problems associated with graphs. A graph  $G = (V, E)$  is a set of  $n$  vertices  $V$  and  $m$  edges  $E \subseteq V \times V$ .  $G$  is said to be undirected if  $(i, j) \in E \Leftrightarrow (j, i) \in E$ , and directed otherwise. We write  $v \rightarrow w$  if there is an edge from  $v$  to  $w$ . Conventionally, vertices and edges in directed graphs are known as nodes and arcs, respectively. One way to specify  $G$  is by its adjacency matrix: an  $n \times n$  matrix  $A$  where  $A_{ij} = 1$  if  $(i, j) \in E$ , and  $A_{ij} = 0$  otherwise. Alternatively,  $G$  can be specified by an adjacency list, which associates with each vertex a list of integers identifying which its neighbours are.

A particularly important graph problem is known as **PATH**. We are given as input a graph  $G$  (in one of the above representations) and two vertices  $s$  and  $t$ . Our task is to decide whether there is a path in  $G$  from  $s$  to  $t$ , i.e. a sequence  $s \rightarrow v_1 \rightarrow v_2 \rightarrow \dots \rightarrow t$ .

We could have defined **PATH** as (for example) the language specified as the subset of  $\{(A, s, t) : A \in \{0,1\}^*, s, t \in \{0,1\}^*\}$  such that  $|A| = n^2$  for some integer  $n$ ,  $1 \leq s, t \leq n$ , and there is a path from  $s$  to  $t$  in the graph corresponding to the adjacency matrix  $A$ , but this level of formality rapidly becomes tedious. In future, when we talk about problems of the form “given  $x \in S$ , determine whether  $x$  satisfies property  $P$ ”, we will assume that  $x$  is specified in a sensible manner, and that it is possible to easily determine whether  $x \in S$ .

## 4.2 Time complexity

We are now ready to define the first complexity class we will consider.

**Definition 4.1.** Let  $T : \mathbb{N} \rightarrow \mathbb{N}$  be a function. A language  $\mathcal{L} \subseteq \{0,1\}^*$  is said to be in  $\text{DTIME}(T(n))$  if there is a multiple tape Turing machine that runs in time  $cT(n)$  for some constant  $c > 0$  and decides  $\mathcal{L}$ .

<sup>1</sup>Later on we will generalise this definition to functional problems.

The Linear Speedup Theorem justifies our choosing arbitrary  $c > 0$  in this definition. The “D” stands for “deterministic”; later on we will also study nondeterministic and randomised models.

By contrast with the Linear Speedup Theorem, it turns out that some problems really do require more time to be solved than others. That is, if we allow a Turing machine more time, it can solve more problems. In what follows we will need to restrict ourselves to so-called time-constructible functions for technical reasons;  $T : \mathbb{N} \rightarrow \mathbb{N}$  is said to be time-constructible if  $T(n) \geq n$  and there is a multiple tape Turing machine  $M$  which computes the function  $M(x) = T(|x|)$ ,  $x \in \{0, 1\}^*$ , in time  $O(T(n))$ . Time-constructibility is not a very significant restriction, as most “natural” functions (e.g. polynomials, exponentials) are time-constructible.

**Theorem 4.2** (Time Hierarchy Theorem, weak version). *If  $f(n)$  is time-constructible,  $\text{DTIME}(f(n))$  is strictly contained within  $\text{DTIME}((f(2n + 1))^3)$ .*

*Proof.* For any time-constructible  $f$ , consider the following language:

$$H_f = \{(\underline{M}, x) : M \text{ accepts } x \text{ in } f(|x|) \text{ steps}\}.$$

Here  $\underline{M}$  is (the binary description of) a multiple tape Turing machine, and  $x$  is its input.  $H_f$  can be decided as follows. First, calculate  $f(|x|)$  and write its binary representation onto a work tape. We will use this as a counter to determine when we need to stop simulating  $M$ . We now run the universal Turing machine  $U$  to simulate  $M$  for  $f(|x|)$  steps, requiring time at most  $O(f(|x|)^3)$ . The construction of  $U$  given in Theorem 3.2 would normally use at most  $O(f(|x|)^2)$  steps for this simulation (because of the quadratic penalty we obtain from simulating  $M$  by a single tape Turing machine), but at each step we need to decrement the counter;  $O(f(|x|))$  time is a generous upper bound to do this. We therefore have  $H_f \in \text{DTIME}(O(f(n)^3))$ . By the Linear Speedup Theorem, this implies  $H_f \in \text{DTIME}(f(n)^3)$ .

On the other hand, we will show that  $H_f \notin \text{DTIME}(f(\lfloor n/2 \rfloor))$  by a diagonalisation argument. Suppose towards a contradiction that there does exist a Turing machine  $M_{H_f}$  which decides  $H_f$  in  $f(\lfloor n/2 \rfloor)$  steps. Then define a machine  $D_f$  which computes

$$D_f(\underline{M}) = \begin{cases} 0 & \text{if } M_{H_f}(\underline{M}, \underline{M}) = 1 \\ 1 & \text{otherwise.} \end{cases}$$

Consider applying  $D_f$  to itself, i.e. computing  $D_f(\underline{D_f})$ , and assume that the answer is 0. This implies that  $M_{H_f}(\underline{D_f}, \underline{D_f}) = 1$  and hence that  $D_f$  accepts  $\underline{D_f}$ , contradicting the assumption. Conversely, assume that the answer is 1. This implies that  $D_f$  does not accept  $\underline{D_f}$  within  $f(\lfloor D_f \rfloor)$  steps. But for any  $M$ ,  $D_f(\underline{M})$  uses the same number of steps as  $M_{H_f}$  does on input  $(\underline{M}, \underline{M})$ , i.e. at most  $f(\lfloor (2|\underline{M}| + 1)/2 \rfloor) = f(\lfloor |\underline{M}| \rfloor)$  steps. So if  $D_f$  accepts  $\underline{D_f}$  it does so within  $f(\lfloor D_f \rfloor)$  steps. We have reached a contradiction.

Combining these two claims, we have  $\text{DTIME}(f(n)) \subset \text{DTIME}((f(2n + 1))^3)$ . □

By using a more efficient universal Turing machine, one can improve this to the following theorem, which we state without proof (see Arora-Barak chapter 1).

**Theorem 4.3** (Time Hierarchy Theorem). *If  $f$  and  $g$  are time-constructible functions such that  $f(n) \log f(n) = o(g(n))$ , then*

$$\text{DTIME}(f(n)) \subset \text{DTIME}(g(n)).$$

### 4.3 Polynomial time

We now come to an important concept, the complexity class  $P$ . This class is simply defined by

$$P = \bigcup_{c \geq 1} \text{DTIME}(n^c).$$

Crucially,  $c$  does not grow with  $n$ . In words,  $P$  is the family of languages which can be decided by a Turing machine in polynomial time (i.e. time which is polynomial in the input size). This class will capture our notion of “efficient computation”. Why is this a reasonable notion of efficiency? After all, an algorithm which runs in time  $\Theta(n^{100})$  can hardly be said to be “efficient”; for some applications, even algorithms which run in time  $\Theta(n^2)$  may be too slow. A simple empirical reason to like the class  $P$  is that in most cases (but not all!) where we have an algorithm for a problem which runs in polynomial time, the degree of the polynomial is in fact reasonable (say 3 or 4).

A class of languages which certainly does not correspond to a real-world notion of efficiency is

$$\text{EXP} = \bigcup_{c \geq 1} \text{DTIME}(2^{n^c}),$$

the class of languages which can be decided in exponential time. In particular,

**Corollary 4.4.**  $P \subset \text{EXP}$ .

*Proof.* As  $n^c = O(2^n)$  for any constant  $c$ ,  $\text{DTIME}(n^c) \subseteq \text{DTIME}(2^n)$ . On the other hand, by the Time Hierarchy Theorem  $\text{DTIME}(2^n) \subset \text{DTIME}((2^{2n+1})^3) \subseteq \text{EXP}$ .  $\square$

To stress the vast gulf between polynomial and exponential time, consider the following thought experiment. Imagine we have two algorithms for some problem, the first of which runs in time  $n^2$ , the second in time  $2^n$ , and we execute these algorithms on a computer which performs one elementary operation per microsecond ( $10^{-6}$  seconds). Then, on a problem instance of size 100, the first algorithm will complete in a hundredth of a second; however, after 40 quadrillion years we will still be waiting for the second algorithm to complete.

There are many problems for which it is easy to find an exponential-time algorithm, but it is far more challenging to find an algorithm which runs in polynomial time. A good example is provided by the problem INTEGER FACTORISATION: given an  $n$ -digit integer  $N$  and an integer  $K$ , both expressed in binary, does  $N$  have a prime factor smaller than  $K$ ? There is an easy exponential-time algorithm for this problem: simply try every possible prime number  $2 \leq j \leq K$ , and see if  $j$  divides  $N$ . The most efficient algorithm currently known for INTEGER FACTORISATION runs in time  $e^{O(n^{1/3}(\log n)^{2/3})}$  and is based on advanced number-theoretic ideas. However, it is not known whether there exists a polynomial-time algorithm for this problem.

An important tool for us to understand which problems are in  $P$  is the notion of polynomial-time reductions. We say that  $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$  is a *polynomial-time computable function* if there is a Turing machine that, when started with input  $x$  on the tape, halts in time  $\text{poly}(|x|)$  with output  $f(x)$ . By analogy with  $P$ ,  $FP$  is defined to be the set of all polynomial-time computable functions. We then say that the language  $\mathcal{A}$  is polynomial time reducible to  $\mathcal{B}$  if there is a polynomial-time computable function  $f$  such that  $w \in \mathcal{A}$  if and only if  $f(w) \in \mathcal{B}$ . Here  $f$  is called a polynomial-time reduction from  $\mathcal{A}$  to  $\mathcal{B}$  and we write  $\mathcal{A} \leq_P \mathcal{B}$ .

It is clear from these definitions that, if  $\mathcal{A} \leq_P \mathcal{B}$  and  $\mathcal{B} \in P$ ,  $\mathcal{A} \in P$ . A nice aspect of polynomial-time reductions is that they compose: if  $\mathcal{A} \leq_P \mathcal{B}$ , and  $\mathcal{B} \leq_P \mathcal{C}$ , then  $\mathcal{A} \leq_P \mathcal{C}$ . This is another major motivation for choosing  $P$  as our class of “reasonable” problems.



## 4.4 A good algorithm: finding paths

There are many problems for which a naïve algorithm would run in exponential time, but which are nevertheless in P. We now give a couple of examples of this phenomenon. From now on, to avoid the painful prospect of calculating time complexities directly in the Turing machine model, we will take a more informal approach and calculate running time based on a model of an idealised “real computer”. For example, we will assume that any symbol in the input can be accessed in time  $O(1)$ . As our main concern will be to distinguish polynomial-time algorithms from exponential-time algorithms, the details of this model are not so important; all that matters is that  $n$  steps of a computation in such a model can be simulated in  $\text{poly}(n)$  time on a Turing machine.

**Theorem 4.5.** *PATH is in P.*

Here the input size is  $O(n^2)$  bits, so we look for an algorithm which runs in time  $\text{poly}(n^2) = \text{poly}(n)$ . Before proving this theorem, we observe that there is a trivial *inefficient* algorithm for PATH: simply enumerate all valid paths of length up to  $n$  that start from vertex  $s$  and see if one of them ends up in  $t$ . However, there could be as many as  $n^n$  such paths, so this algorithm uses time super-polynomial in  $n$ .

*Proof of Theorem 4.5.* Recall that in the PATH problem we are given a directed graph  $G = (V, E)$  and two nodes  $s$  and  $t$ , and have to decide if there is a path from  $s$  to  $t$ . Our algorithm will be based on a data structure known as a queue, which allows the operations of addition (to the end of the queue) and removal (from the start of the queue). The algorithm proceeds as follows.

1. Mark node  $s$  and add it to the queue.
2. While the queue is not empty:
  - (a) Remove the first node  $y$  from the queue.
  - (b) For all edges  $(y, u)$  such that  $u$  is not marked, mark  $u$  and add  $u$  to the end of the queue.
3. Accept if  $t$  is marked.

It is easy to see that this algorithm is correct; we now discuss its running time. A queue can be implemented such that enqueueing and dequeuing an element takes only time  $O(1)$ . Each node is only added to the queue at most once. Therefore, an upper bound on the running time is given by the sum over all nodes  $y$  of the time required to check each of  $y$ 's neighbours. If the input graph  $G$  has  $n$  vertices and  $m$  edges and is given in terms of an adjacency matrix, we get a bound of  $O(n^2)$ , and if  $G$  is given as an adjacency list we get a bound of only  $O(n + m)$ , as each edge is checked only once<sup>1</sup>. In either case, this is polynomial in the input size.  $\square$

The above algorithm is known as breadth-first search. Observe that the algorithm can easily be modified to actually output a path from  $s$  to  $t$  (if one exists), by maintaining a tree containing all nodes visited thus far. When we mark a node  $u$  neighbouring a node  $y$ , we add it to the tree as a leaf descending from  $y$ . Then, if we find  $t$ , we explore the tree in reverse, repeatedly visiting the parent of the current node until we get back to  $s$ . Reversing the direction of each of these arcs gives a path from  $s$  to  $t$ .

---

<sup>1</sup>The meaning of big-O notation may not be clear in the case of expressions involving more than one independently growing parameter. Here, we tacitly understand  $m$  to be a function of  $n$ .

		<i>a</i>	<i>b</i>	<i>b</i>	<i>a</i>
	0	1	2	3	4
<i>b</i>	1	1	1	2	3
<i>a</i>	2	1	2	2	2
<i>c</i>	3	2	2	3	3
<i>a</i>	4	3	3	3	3

Figure 5: A completed table demonstrating that  $\text{ED}(baca, abba) = 3$ .

#### 4.5 Another good algorithm: computing edit distance

The second example we will consider is the problem of calculating edit distance between strings, an important problem in bioinformatics and elsewhere. Given two strings  $x, y \in \Sigma^*$ , the edit distance  $\text{ED}(x, y)$  is the number of edits required to change  $x$  into  $y$ , where an edit consists of removing, inserting or replacing a symbol in  $x$ . For example,  $\text{ED}(baca, abba) = 3$  via the sequence of edits

$$baca \mapsto baba \mapsto aba \mapsto abba.$$

For any given pair of strings there are infinitely many sequences of edits which can transform one into the other. However, it turns out that we can find a shortest such sequence very efficiently.

**Theorem 4.6.** *Given two strings  $x, y$  of length at most  $n$ ,  $\text{ED}(x, y)$  can be computed in time  $O(n^2)$ .*

*Proof.* The technique we will use is known as dynamic programming, which is an important tool in algorithm design. Roughly speaking, the method works by breaking a problem down into simpler overlapping subproblems, then efficiently combining the solutions to the subproblems. Assume the string  $x$  is of length  $m$ , and  $y$  is of length  $n$ . Let  $x[i]$  denote the string consisting of the first  $i$  characters of  $x$  (if  $i = 0$ ,  $x[i]$  is the empty string). We construct a  $(m + 1) \times (n + 1)$  table  $E$  such that  $E_{ij} = \text{ED}(x[i], y[j])$  (where  $0 \leq i \leq m$  and  $0 \leq j \leq n$ ). This can be done using the observation (which should be checked!) that

$$E_{ij} = \begin{cases} i & \text{if } j = 0 \\ j & \text{if } i = 0 \\ \min\{E_{i-1,j} + 1, E_{i,j-1} + 1, E_{i-1,j-1} + [x_i \neq y_j]\} & \text{otherwise.} \end{cases}$$

Here we define  $[x_i \neq y_j] = 1$  if  $x_i \neq y_j$  and  $[x_i \neq y_j] = 0$  if  $x_i = y_j$ . Thus we can first fill in the first row and column of  $E$  (which do not depend on  $x$  and  $y$ ), then fill the remaining entries of  $E$  row by row, starting at the top left.  $\text{ED}(x, y)$  is then the bottom-right entry in  $E$ . As there are  $(m + 1)(n + 1)$  entries in  $E$ , and each requires  $O(1)$  time to compute, the time required to compute  $\text{ED}(x, y)$  is  $O(mn)$ . This process is illustrated in Figure 5.  $\square$

#### 4.6 One more good algorithm: finding a maximum matching

The final problem we will consider is graph-theoretic in nature: finding a maximum matching in a bipartite graph. A *matching*  $M$  in an undirected graph  $G$  is a subset of the edges of  $G$  such that no pair of edges has a vertex in common.  $M$  is said to be a *maximum* matching if it is as large




Figure 6: A bipartite graph with a (non-maximum) matching, and an augmenting path with respect to that matching, which results in a perfect matching.

as possible, and *perfect* if every vertex in  $G$  is included in an edge of  $M$ . Finally,  $G$  is said to be bipartite if its vertices can be partitioned into sets  $L$  and  $R$  such that every edge in  $G$  connects a vertex in  $L$  to a vertex in  $R$ . See Figure 6 for an illustration.

If we let the sets  $L$  and  $R$  correspond to “workers” and “jobs”, and an edge correspond to “this worker can do that job”,  $G$  has a perfect matching if and only if every job can be assigned to a worker. A simple criterion is known for whether such a perfect matching exists.

**Theorem 4.7** (Hall’s marriage theorem). *Let  $G = (V, E)$  be a bipartite undirected graph with bipartition  $V = L \cup R$ , and  $|L| = |R|$ . Then  $G$  has a perfect matching if and only if, for all  $X \subseteq L$ ,  $|X| \leq |\{(x, y) : x \in X, (x, y) \in E\}|$ .*

Observe that Hall’s marriage theorem does not imply an efficient algorithm to determine whether  $G$  has a perfect matching. If  $|L| = |R| = n$ , to verify the conditions of the theorem could require testing  $2^n$  different subsets of  $L$ . However, a different approach does lead to an efficient algorithm; in fact, an efficient algorithm for the more general task of *finding* a maximum matching.

The approach we will use is known as the augmenting path algorithm. Let  $M \subseteq E$  be a matching in a bipartite graph  $G = (V, E)$ . We say that  $v \in V$  is unmatched if  $v$  is not included in  $M$ . An *augmenting path* for  $M$  is a sequence of edges that starts and ends with an unmatched vertex and alternates between edges of  $E \setminus M$  and edges of  $M$ . The algorithm proceeds as follows.

1. Set  $M$  to  $\emptyset$ .
2. While there is an augmenting path  $P$  for  $M$ , replace  $M$  with  $M \Delta P$ .
3. If there is no augmenting path, return  $M$ .

From the definition of an augmenting path, we see that replacing  $M$  with  $M \Delta P$  increases the size of  $M$  by 1; a little thought shows that the resulting set is still a matching. Augmenting paths have the following additional nice property.

**Lemma 4.8.** *Let  $M$  be a matching. If  $M$  is not maximum, there exists an augmenting path for  $M$ .*

*Proof.* Assume that  $M$  is not maximum, let  $N$  be a maximum matching, and consider the graph with edges  $X = M \Delta N$ . This graph is of degree at most 2, so consists of a collection of disconnected paths and cycles, where a cycle is a path whose initial vertex is the same as its final vertex, i.e. a path of the form  $v_1 \rightarrow v_2 \rightarrow v_3 \rightarrow \dots \rightarrow v_1$ . Each path and cycle alternates between edges from  $M$  and edges from  $N$ . As  $|N| > |M|$ ,  $X$  contains more edges from  $N$  than from  $M$ . Each cycle in  $X$

has the same number of edges from each, so there must exist a path with exactly one more edge from  $N$  than from  $M$ . This is an augmenting path as desired.  $\square$

This lemma implies that the augmenting path algorithm is correct (i.e. will always output a maximum matching of  $M$ ). The remaining question is how to find such an augmenting path efficiently, which can be achieved as follows for bipartite graphs  $G$ . Form a directed graph from  $G$  by directing edges from  $L$  to  $R$  if they do not belong to  $M$ , and from  $R$  to  $L$  if they do, and call this directed graph  $G'$ . Then any path in  $G'$  from an unmatched node in  $L$  to an unmatched node in  $R$  corresponds to an augmenting path in  $G$  with respect to  $M$ . In order to find such a path, we add a new node  $s$  to  $G'$ , and add an arc from  $s$  to every unmatched node in  $L$ . We then apply the algorithm for PATH, starting at  $s$  and terminating when we find an unmatched node in  $R$ . The path which is output gives our desired augmenting path. The whole algorithm runs in time  $\text{poly}(n)$ .

## 4.7 Historical notes and further reading

Papadimitriou chapter 2 has a careful discussion of how a simple model of a “real” computer, known as the RAM model, can be simulated efficiently by a Turing machine.

The Time Hierarchy Theorem is one of the founding results of the modern era of computational complexity, and was proven by Richard Stearns and Juris Hartmanis in 1965.

It should hopefully be clear from the three examples given above that designing efficient algorithms is far from a trivial process. The book CLRS is a comprehensive introduction to efficient algorithms, and includes considerably more detail about breadth-first search and other graph search techniques. Some beautiful notes about dynamic programming have been written by Jeff Erickson<sup>1</sup>. There are also some nice notes about matchings from a course by Uri Zwick<sup>2</sup>. There is a generalisation of the algorithm for finding a maximum matching to arbitrary graphs. This is due to Jack Edmonds, whose 1965 paper called “Paths, trees, and flowers” contains one of the first discussions of the role of efficiency in algorithms.

---

<sup>1</sup><http://cs.uiuc.edu/~jeffe/teaching/algorithms/2009/notes/03-dynprog.pdf>

<sup>2</sup><http://www.cs.tau.ac.il/~zwick/grad-algo-0910/match.pdf>

## 5 Certificates and the class NP

There are many problems in life which we may not necessarily know how to solve, but for which we can verify a claimed solution. This concept motivates the definition of the complexity class NP.

**Definition 5.1.** NP is the class of languages  $\mathcal{L} \subseteq \{0,1\}^*$  such that there exists a polynomial  $p : \mathbb{N} \rightarrow \mathbb{N}$  and a polynomial-time Turing machine  $M$  such that, for all  $x \in \{0,1\}^*$ ,  $x \in \mathcal{L}$  if and only if there exists  $w \in \{0,1\}^{p(|x|)}$  such that  $M(x, w) = 1$ .

We call  $M$  a verifier for  $\mathcal{L}$  and  $w$  a certificate or witness for  $x$ . To illustrate this definition, consider the following examples of languages / decision problems in NP.

- The PATH problem is in NP. Indeed, any language  $\mathcal{L} \in \mathbf{P}$  is automatically in NP; the verifier can simply ignore the certificate  $w$  and decide membership in  $\mathcal{L}$  in polynomial time.
- The language COMPOSITES, defined to be the set of integers  $\{N : N = p \times q, p, q \geq 2\}$ , where  $N$  is expressed in binary. A certificate is the factorisation  $(p, q)$ ; given  $p$  and  $q$ , it can be checked in time polynomial in the input size (i.e.  $\text{poly}(\log N)$ ) that  $N = pq$ . In fact, COMPOSITES is in P but this is not obvious!
- The problem INTEGER FACTORISATION: given two integers  $N, k$ , does  $N$  have a prime factor smaller than  $k$ ? Once again, a certificate is the factorisation of  $N$ . However, in this case INTEGER FACTORISATION is not known to be in P.

Another, and very important, example of a problem in NP is known as boolean satisfiability (SAT). In order to discuss this, we pause to introduce some notation and concepts relating to boolean functions.

### 5.1 Boolean functions

- A boolean variable  $x_i$  takes values 0 or 1 (corresponding to “false” or “true”).
- A boolean function  $f : \{0,1\}^n \rightarrow \{0,1\}$  is a function of  $n$  boolean variables. The list of the values  $f$  takes on each input, written out in lexicographic order, is called the truth table of  $f$ .
- Boolean operations AND ( $x_1 \wedge x_2$ ), OR ( $x_1 \vee x_2$ ), and NOT ( $\neg x$ ) are defined as one might expect:  $x_1 \wedge x_2 = 1$  if and only if  $x_1 = x_2 = 1$ ;  $x_1 \vee x_2 = 1$  if either  $x_1, x_2$  or both are 1;  $\neg x = 1 - x$ . We will also use the operations XOR (addition modulo 2, or  $\oplus$ ) and logical implication ( $\Rightarrow$ ); we summarise all of these below.

$x$	$y$	$x \wedge y$	$x$	$y$	$x \vee y$	$x$	$y$	$x \oplus y$	$x$	$y$	$x \Rightarrow y$
0	0	0	0	0	0	0	0	0	0	0	1
0	1	0	0	1	1	0	1	1	0	1	1
1	0	0	1	0	1	1	0	1	1	0	0
1	1	1	1	1	1	1	1	0	1	1	1

- De Morgan’s laws state that

$$\begin{aligned} \neg(x_1 \wedge x_2 \wedge \cdots \wedge x_k) &= (\neg x_1) \vee (\neg x_2) \vee \cdots \vee (\neg x_k), \text{ and} \\ \neg(x_1 \vee x_2 \vee \cdots \vee x_k) &= (\neg x_1) \wedge (\neg x_2) \wedge \cdots \wedge (\neg x_k). \end{aligned}$$

- A literal is a boolean variable, possibly preceded by a NOT; for example,  $x_1$  and  $\neg x_1$  are literals.
- A boolean formula is an expression containing boolean variables and AND, OR and NOT operations, such as  $\phi(x_1, x_2, x_3) = x_1 \wedge (\neg(\neg x_2 \vee x_3) \wedge x_3)$ .
- A boolean formula is in Conjunctive Normal Form (CNF) if it is written as  $c_1 \wedge c_2 \wedge \dots \wedge c_n$ , where each  $c_i$  is a *clause* which is the OR of one or more literals (e.g. the formula

$$\phi(x_1, x_2, x_3, x_4) = (x_1 \vee \neg x_2) \wedge (\neg x_2 \vee \neg x_3 \vee x_4) \wedge (x_1)$$

is in CNF). Similarly, a formula is in Disjunctive Normal Form (DNF) if it is the OR of clauses, each of which is the AND of one or more literals.

- A boolean formula is *satisfiable* if there is an assignment to the variables that makes it evaluate to true. For example, the previous formula is satisfiable (e.g. set  $x_1 = 1, x_2 = 0, x_3 = 1, x_4 = 0$ ).

We have the following “universality” result for boolean formulae.

**Theorem 5.1.** *Any boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  can be expressed as a boolean formula  $\phi_f$  in CNF.*

*Proof.* For any boolean function  $f$ , we can generate a boolean formula  $\phi_f$  which represents  $f$  as follows. For each input  $y \in \{0, 1\}^n$  such that  $f(y) = 1$ , write down the clause  $c_y = (\ell_1 \wedge \ell_2 \wedge \dots \wedge \ell_n)$ , where  $\ell_i = x_i$  if  $y_i = 1$ , and  $\ell_i = \neg x_i$  otherwise. Then set  $\phi_f(x) = \bigvee_{y, f(y)=1} c_y(x)$ . As  $c_y(x)$  evaluates to true if and only if  $x = y$ , it is clear that, for all  $x$ ,  $\phi_f(x) = f(x)$ . This procedure produces a formula in DNF. To obtain a CNF formula, take the negation of the DNF formula  $\phi_{\neg f}$  and use De Morgan’s law.  $\square$

We are now equipped to define SAT, which is the following problem: **Given a boolean formula  $\phi$  in CNF, is  $\phi$  satisfiable?**

To check that  $\text{SAT} \in \text{NP}$ , observe that if we are given an assignment to the variables of  $\phi$  which is claimed to make  $\phi$  evaluate to true, we can just plug it in and check. However, it is not clear how to (efficiently) *find* an assignment that makes  $\phi$  true. One can iterate over all possible assignments to the variables of  $\phi$ ; however, if  $\phi$  depends on  $n$  variables, this takes time  $\Omega(2^n)$ , which is exponential in the input size (assuming that the number of clauses is  $O(n)$ ).

## 5.2 Relationship between NP and other classes

We have the following straightforward containments.

**Theorem 5.2.**  $\text{P} \subseteq \text{NP} \subseteq \text{EXP}$ .

*Proof.* For the first inclusion, suppose that  $\mathcal{L} \in \text{P}$ , so there exists a polynomial-time Turing machine  $N$  which decides  $\mathcal{L}$ . Then, given an empty certificate  $w$ , the verifier can simply use  $N$  to decide  $\mathcal{L}$ . For the second inclusion, if  $\mathcal{L} \in \text{NP}$ , we can decide  $\mathcal{L}$  in time  $2^{O(p(n))}$  by trying all possible certificates  $w \in \{0, 1\}^{p(n)}$  as inputs to the verifier  $M$  in turn. As  $p(n) = O(n^c)$  for some  $c > 1$ ,  $\mathcal{L} \in \text{EXP}$ .  $\square$

Perhaps surprisingly, for some problems in NP, no algorithm is known that runs significantly faster than this naïve exponential-time algorithm.

### 5.3 NP-completeness

The concept of NP-completeness is a way of formalising the idea of what the “hardest” problems in NP are, and is defined as follows.

- We say that a language  $\mathcal{A}$  is NP-hard if, for all  $\mathcal{B} \in \text{NP}$ ,  $\mathcal{B} \leq_P \mathcal{A}$ .
- We say that a language  $\mathcal{A}$  is NP-complete if  $\mathcal{A} \in \text{NP}$ , and  $\mathcal{A}$  is NP-hard.

So, if  $\mathcal{A}$  is NP-hard, deciding membership in  $\mathcal{A}$  is, up to polynomial terms in the running time, at least as hard as the hardest problem in NP. This can be interpreted as evidence that there is no polynomial-time algorithm for  $\mathcal{A}$ , because if this were true, there would be a polynomial-time algorithm for every problem in NP, and we would have  $P = NP$ . The probability of this being the case is debatable, but it is generally considered very unlikely. In the next section we will see some substantial evidence that  $P \neq NP$ , at least if one is a pessimist; the evidence being that a number of important problems turn out to be NP-complete.

However, it is perhaps not obvious that there should exist *any* NP-complete problems, let alone any “natural” ones. The remarkable fact that such problems do exist is called the Cook-Levin Theorem.

**Theorem 5.3** (Cook-Levin Theorem). *SAT is NP-complete.*

In order to prove this theorem, we will need to understand another, equivalent, definition of the class NP, as the class of languages recognised by a *nondeterministic* Turing machine. In fact, this was the original way in which NP was defined, and NP stands for “nondeterministic polynomial time” (**not** “non-polynomial time”).

### 5.4 Nondeterministic Turing machines

Nondeterministic Turing machines (NDTMs) are (nonphysical and unrealistic!) generalisations of the Turing machine model. Instead of having one transition function  $\delta$ , an NDTM may have several functions  $\delta_1, \dots, \delta_K$ . A computational path is a sequence of choices of transition functions (i.e. integers between 1 and  $K$ ). We think of an NDTM  $M$  as making all of these transitions in parallel, leading to many different potential computational paths. See Figure 7 for an illustration of this. NDTMs also have a special ACCEPT state.

We say that an NDTM  $M$  decides a language  $\mathcal{L} \subseteq \{0, 1\}^*$  if:

- All  $M$ 's computational paths reach either the state HALT or the state ACCEPT;
- For all  $x \in \mathcal{L}$ , on input  $x$  at least one path reaches the ACCEPT state;
- For all  $x \notin \mathcal{L}$ , on input  $x$  no path reaches the ACCEPT state.

We say that  $M$  runs in time  $T(n)$  if, for every input  $x \in \{0, 1\}^*$  and every sequence of nondeterministic choices,  $M$  reaches either the state HALT or the state ACCEPT in at most  $T(|x|)$  steps. We can now define the class NTIME as follows, by analogy with DTIME:

**Definition 5.2.** A language  $\mathcal{L} \subseteq \{0, 1\}^*$  is said to be in  $\text{NTIME}(T(n))$  if there is an NDTM that runs in time  $cT(n)$  for some constant  $c > 0$  and decides  $\mathcal{L}$ .




Figure 7: An NDTM can take multiple paths “in parallel”.

Unlike standard Turing machines, we do not expect to be able to actually implement NDTMs. However, they will nevertheless be a very useful conceptual tool.

We first observe that there is no loss of generality in assuming that  $K = 2$ , i.e. that  $M$  has only two possible transition rules. Indeed, for any  $K \geq 2$  one can associate each transition rule  $\delta$  with a string of  $\lceil \log_2 K \rceil$  bits and create new rules  $\delta'_0, \delta'_1$  which, over  $\lceil \log_2 K \rceil$  steps, select the bits of  $\delta$  (storing each bit in a different head state).

**Theorem 5.4.**  $\text{NP} = \bigcup_{c \geq 1} \text{NTIME}(n^c)$ .

*Proof.* First, suppose  $\mathcal{L} \in \text{NP}$ . Then there exists a polynomial-time verifier  $V$  such that, for all  $x \in \mathcal{L}$ , there is a certificate  $w$  of size  $\text{poly}(|x|)$  such that  $V$  accepts on input  $(x, w)$ ; and for all  $x \notin \mathcal{L}$ , no certificate  $w$  exists such that  $V$  accepts on input  $(x, w)$ . To decide  $\mathcal{L}$ , on input  $x$  our NDTM  $M$  simply guesses  $w$  nondeterministically and runs  $V$  on  $(x, w)$ . To make this idea of “guessing” concrete, we imagine that  $M$  has two special transition functions  $\delta_0, \delta_1$ , which correspond to writing a 0 or a 1 to the tape, allowing  $M$  to guess each bit of  $w$  in turn and write it to the tape.

Second, suppose  $M$  is a polynomial-time NDTM that decides  $\mathcal{L}$ . Then, for each  $x \in \mathcal{L}$ , there is a computational path of length  $\text{poly}(|x|)$  leading to the ACCEPT state, but for all  $x \notin \mathcal{L}$ , there is no such path. This serves as a certificate: the verifier takes as input  $(x, p)$ , where  $p$  identifies a computational path, and just simulates the path  $p$  on input  $x$ .  $\square$

One can similarly define the nondeterministic analogue of EXP,  $\text{NEXP} = \bigcup_{c \geq 1} \text{NTIME}(2^{n^c})$ .

## 5.5 Proof of the Cook-Levin theorem

We already know that  $\text{SAT} \in \text{NP}$ . To show that SAT is in fact NP-complete, we need to show that every language in NP reduces to SAT. To be precise, we will show that, for any language  $\mathcal{L} \subseteq \{0, 1\}^*$  such that  $\mathcal{L} \in \text{NP}$ , given  $x \in \{0, 1\}^n$ , we can construct (in time  $\text{poly}(n)$ ) a CNF formula  $\phi$  such that  $\phi$  is satisfiable if and only if  $x \in \mathcal{L}$ . As  $\mathcal{L} \in \text{NP}$ , there is a polynomial-time NDTM  $M$  that decides  $\mathcal{L}$ . Given a description of  $M$ , we will encode this as a formula  $\phi$ , which will evaluate to true if and only if  $M$  has an accepting path on  $x$ .

As  $M$  is a polynomial-time NDTM, there is a constant  $c$  such that  $M$  runs in time at most  $T = n^c$  on input  $x$ . We associate a  $(T + 1) \times (T + 2)$  tableau with each computational path of  $M$ , where row  $t$  contains the configuration of  $M$  at step  $t$  of the path; see Figure 8 for an illustration.



▷	START	$x_1$	$x_2$	...	$x_n$	□	...	□
▷	0	$s$	$x_2$	...	$x_n$	□	...	□
⋮								
▷	0	1	ACCEPT	...	1	0	...	□

Figure 8: A tableau describing a particular computational path.

Each row thus consists of a triple  $(\ell, q, r)$ , where  $q$  is a head state and  $\ell$  and  $r$  are strings of tape symbols. Then our CNF formula is made up of subformulae,

$$\phi = \phi_{\text{cell}} \wedge \phi_{\text{start}} \wedge \phi_{\text{move}} \wedge \phi_{\text{accept}},$$

where

- $\phi_{\text{cell}}$  evaluates to true if all squares in the tableau are uniquely filled;
- $\phi_{\text{start}}$  evaluates to true if the first row is the correct starting configuration;
- $\phi_{\text{move}}$  evaluates to true if the tableau is a valid computational path of  $M$  (according to its transition rules);
- $\phi_{\text{accept}}$  evaluates to true if the tableau contains an accepting state.

It should be intuitively clear that there is a satisfying assignment to  $\phi$  if and only if there is a valid tableau (i.e. computational path that accepts  $x$ ). Formally, introduce a set of boolean variables  $c_{i,j,s}$ , each of which evaluates to true if cell  $(i, j)$  contains symbol  $s$ . Then

$$\begin{aligned} \phi_{\text{accept}} &= \bigvee_{i,j} c_{i,j,\text{ACCEPT}}; \\ \phi_{\text{start}} &= c_{1,1,\triangleright} \wedge c_{1,2,\text{START}} \wedge c_{1,3,x_1} \wedge c_{1,4,x_2} \wedge \cdots \wedge c_{1,n+2,x_n} \wedge c_{1,n+3,\square} \wedge \cdots \wedge c_{1,T+2,\square}; \\ \phi_{\text{cell}} &= \bigwedge_{i,j} \left( \left( \bigvee_s c_{i,j,s} \right) \bigwedge_{t \neq u} (\neg c_{i,j,t} \vee \neg c_{i,j,u}) \right). \end{aligned}$$

The last of these encodes the constraint that every cell has a value, and no cell has two values, and can easily be written in CNF. For the final constraint  $\phi_{\text{move}}$ , we need to encode the validity of a computational path. The key insight is that this can be done using the *locality* of Turing machines. Consider a  $2 \times 3$  “window” (submatrix) in a tableau. If  $S$  is the number of possible symbols used to write a configuration, there are only  $\text{poly}(S)$  possible windows – a fixed, finite number. However, some of these are disallowed (“illegal”) because they correspond to transitions which cannot take place. For example, the head cannot move two positions in one step, so if  $s$  is a head state the window

$s$	0	0
0	0	$s$

is illegal. In general, whether a given window is legal or not will depend on the transition functions of  $M$ . The constraint that a window  $w$  is legal can be written as

$$\phi_w = \bigvee_{\text{legal windows } v} [w_{ij} = v_{ij} \text{ for } i \in \{1, 2\}, j \in \{1, 2, 3\}].$$

Note that this constraint is a boolean formula with a fixed, finite number of terms, which by Theorem 5.1 can be written in CNF.

We claim that, if all windows in a tableau are legal, then each row in the tableau corresponds to a configuration which legally follows the previous one, i.e. the tableau as a whole corresponds to a valid computational path. Intuitively, this is because the head can only move one step at a time, so every possible transition will occur within some window. The (somewhat more formal) proof is by induction. For the base case,  $\phi_{\text{start}}$  ensures that the first row is valid. We now show that, if row  $r$  is a valid configuration, then row  $r + 1$  is also a valid configuration. Consider any window onto rows  $r$  and  $r + 1$ . If the window does not contain a head symbol in its top row, in order for the window to be legal the middle symbol in its bottom row must be equal to the middle symbol in its top row, so all tape cells which are not adjacent to the head are preserved<sup>1</sup>. On the other hand, as row  $r$  is a valid configuration, there must exist a window containing a head symbol in the middle of its top row. This window will only be legal if its bottom row corresponds to a valid transition (and in particular also contains a head symbol). And as the head can only move at most one position per step, a window whose first row does not contain a head symbol cannot contain a head symbol in the middle of its bottom row, so row  $r + 1$  must contain exactly one head symbol. Thus row  $r + 1$  is a valid configuration.

We can therefore write

$$\phi_{\text{move}} = \bigwedge_{\text{windows } w} \phi_w.$$

Thus, combining the four subformulae,  $\phi$  being satisfiable is equivalent to  $M$  accepting  $x$ . Finally, observe that  $\phi$  is of size  $\text{poly}(T) = \text{poly}(n)$ , and can be produced in time  $\text{poly}(n)$  from a description of  $M$ . This completes the proof.

## 5.6 Historical notes and further reading

The Cook-Levin Theorem was proven by Stephen Cook in 1971, and independently by Leonid Levin in 1973, on the other side of the Iron Curtain. Every good textbook on computational complexity includes a proof of this fundamental theorem. The proof here is largely based on that in Sipser (chapter 7).

A poll from about a decade ago gives a good snapshot of the complexity theory community's opinions on the P vs. NP question<sup>2</sup>. Scott Aaronson has a nice list of non-rigorous arguments to believe that  $P \neq \text{NP}$ <sup>3</sup>.

---

<sup>1</sup>The alert reader may wonder about the special cases of the first and last columns of the tableau. These can be dealt with either by having special constraints on whether windows involving these columns are legal, or introducing special “border” symbols to identify the start and end of each row.

<sup>2</sup><http://www.cs.umd.edu/~gasarch/papers/poll.pdf>

<sup>3</sup><http://www.scottaaronson.com/blog/?p=122>




Figure 9: Examples of a Hamiltonian path in an undirected graph and a 3-colouring.

## 6 Some more NP-complete problems

It turns out that many, many interesting problems are known to be NP-complete. Garey and Johnson alone list over 300! The Cook-Levin theorem allows us to prove a given problem to be NP-complete in a much more straightforward manner than was the case for SAT. Indeed, if we can show, for some language  $\mathcal{L} \in \text{NP}$ , that  $\text{SAT} \leq_P \mathcal{L}$ , then  $\mathcal{L}$  is NP-complete. Here are a few examples of other problems which are known to be NP-complete.

- **CLIQUE**: given an undirected graph  $G$  and an integer  $k$ , does  $G$  contain a clique on  $k$  vertices, i.e. a  $k$ -subset of the vertices such that every pair of vertices in the subset are connected by an edge?
- **HAMILTONIAN PATH**: given a directed graph  $G$ , does it contain a path visiting each node exactly once?
- **SUBSET SUM**: given a sequence  $S$  of  $n$  integers and a “target”  $t$ , is there a subsequence of  $S$  that sums to  $t$ ?
- **QUADRATIC DIOPHANTINE EQUATION**: given positive integers  $a$ ,  $b$  and  $c$ , are there positive integers  $x$  and  $y$  such that  $ax^2 + by = c$ ?
- **SHORTEST COMMON SUPERSTRING**: given a finite set of strings  $S \subset \{0, 1\}^*$  and an integer  $k$ , is there a string  $s \in \{0, 1\}^*$  such that  $|s| \leq k$  and each  $x \in S$  is a substring of  $s$ ?
- **3-COLOURING**: given a graph  $G$ , can the vertices each be assigned one of three colours, such that adjacent vertices receive different colours?

We will prove NP-completeness for the first three of these problems. However, the first NP-completeness proof we will give is for a restricted variant of SAT, which turns out to be very useful for other such proofs. We say that  $\phi$  is a  $k$ -CNF formula if it is a boolean formula in CNF with at most  $k$  variables per clause. Let  $k$ -SAT denote the special case of the SAT problem where the input is restricted to  $k$ -CNF formulae. For example,

$$\phi = (x_1 \vee \neg x_2) \wedge (\neg x_1 \vee x_5 \vee x_7) \wedge (x_3)$$

is a valid 3-SAT instance.

**Theorem 6.1.** *3-SAT is NP-complete.*

*Proof.* As 3-SAT is a restriction of SAT, it is clearly in NP. To show that it is NP-hard, we reduce SAT to 3-SAT. Given as input a boolean formula  $\phi$  in CNF, we will obtain a new formula  $\phi'$  by replacing each clause which contains  $k \geq 4$  literals  $(\ell_1, \dots, \ell_k)$  with  $k - 2$  clauses which contain 3 literals, while preserving satisfiability of  $\phi$ . Consider the formula

$$\phi_{\ell_1, \dots, \ell_k} = (\ell_1 \vee \ell_2 \vee u_1) \wedge (\ell_3 \vee \neg u_1 \vee u_2) \wedge \dots \wedge (\ell_{k-2} \vee \neg u_{k-4} \vee u_{k-3}) \wedge (\ell_{k-1} \vee \ell_k \vee \neg u_{k-3}),$$

where we have introduced  $k - 3$  new variables  $u_1, \dots, u_{k-3}$ . If any of the original literals  $\ell_i$  evaluates to true, then the whole formula is satisfiable (taking  $u_j = 1$  for  $u_j$  up to the clause in which  $\ell_i$  is encountered, and  $u_j = 0$  thereafter). On the other hand, if none of the  $\ell_i$  evaluate to true,  $\phi_{\ell_1, \dots, \ell_k}$  is equivalent to

$$u_1 \wedge (\neg u_1 \vee u_2) \wedge \dots \wedge (\neg u_{k-4} \vee u_{k-3}) \wedge (\neg u_{k-3}),$$

and a moment's thought shows that this is not satisfiable. Thus the resulting overall formula  $\phi'$  obtained by replacing all clauses containing more than 3 literals in this way is satisfiable if and only if  $\phi$  was satisfiable, so  $\text{SAT} \leq_P \text{3-SAT}$ .  $\square$

Can we improve this result to show that 2-SAT is NP-complete? Perhaps surprisingly, we have the following result suggesting that we cannot.

**Theorem 6.2.** *2-SAT*  $\in$  P.

*Proof.* We will reduce 2-SAT to PATH. Given an  $n$ -variable 2-SAT formula  $\phi$  with exactly two distinct variables per clause (clauses containing only one variable can be dealt with trivially), we construct a directed graph  $G$  as follows.

- $G$  has  $2n$  nodes labelled by literals (i.e. variables and their negations) appearing in  $\phi$ .
- For each clause  $(\ell_1 \vee \ell_2)$ ,  $G$  has arcs  $\neg \ell_1 \rightarrow \ell_2$  and  $\neg \ell_2 \rightarrow \ell_1$ .

These arcs are intended to symbolise logical implication. If we have a clause  $(x_1 \vee \neg x_2)$ , for example, we can think of it as capturing the implication  $x_2 \Rightarrow x_1$  (“if  $x_2 = 1$ , then  $x_1 = 1$ ”) or equivalently the contrapositive  $\neg x_1 \Rightarrow \neg x_2$ .  $\phi$  is satisfiable if and only if all of these implications are consistent.

We now show that  $\phi$  is unsatisfiable if and only if there is a variable  $x_i$  such that there is a path from  $x_i$  to  $\neg x_i$  and from  $\neg x_i$  to  $x_i$ .

- ( $\Leftarrow$ ): Suppose for a contradiction that there is a path from  $x_i$  to  $\neg x_i$  and from  $\neg x_i$  to  $x_i$ , and yet  $\phi$  is satisfiable. Let  $T$  be a satisfying assignment to the variables. If  $T(x_i) = 1$ , then  $T(\neg x_i) = 0$ , so there must be an arc  $\ell_1 \rightarrow \ell_2$  on the path from  $x_i$  to  $\neg x_i$  such that  $T(\ell_1) = 1$ ,  $T(\ell_2) = 0$ . This means that  $(\neg \ell_1 \vee \ell_2)$  is a clause of  $\phi$ , but is not satisfied by  $T$ . The case  $T(x_i) = 0$  is similar.
- ( $\Rightarrow$ ): We consider the following method for finding a satisfying assignment to  $\phi$  by assigning truth values to each node of  $G$ . Repeatedly, for each node  $\ell$  which has not yet been assigned a value, and such that there is no path from  $\ell$  to  $\neg \ell$ , we assign 1 to  $\ell$  and each node reachable from  $\ell$ ; we also assign 0 to  $\neg \ell$  and the negations of the nodes reachable from  $\ell$  (i.e. the nodes from which  $\neg \ell$  is reachable).

We first check that this process makes sense (i.e. assigns consistent values to the literals). If there were a path from  $\ell$  to two nodes  $m$  and  $\neg m$ , then by symmetry of  $G$  there would be paths from both  $m$  and  $\neg m$  to  $\neg \ell$ , so there would be a path from  $\ell$  to  $\neg \ell$ , contradicting the

hypothesis. So, in each step, truth values are assigned consistently. Further, if there is a path from  $\ell$  to a node  $m$  previously assigned 0 (say), as  $m$  is reachable from  $\ell$ ,  $\ell$  must have also previously been assigned 0.

As we have assumed that for each variable  $x_i$ , there is either no path from  $x_i$  to  $\neg x_i$ , or no path from  $\neg x_i$  to  $x_i$ , all  $x_i$  will be assigned a value. As the assignment is produced such that whenever  $\ell \rightarrow m$ , either  $m$  is assigned 1 or  $\ell$  is assigned 0, the assignment satisfies  $\phi$ .

We can therefore apply the polynomial-time algorithm for PATH given in Theorem 4.5 to each pair  $(x_i, \neg x_i)$  to solve 2-SAT in polynomial time.  $\square$

The next problem we consider is a basic question from arithmetic, the SUBSET SUM problem<sup>1</sup>: given a sequence  $S$  of  $n$  integers and a “target”  $t$ , is there a subsequence of  $S$  that sums to  $t$ ?

**Theorem 6.3.** SUBSET SUM is NP-complete.

*Proof.* It is obvious that SUBSET SUM is in NP (the certificate is simply the subsequence of  $S$  that sums to the target). To prove that it is NP-complete, we reduce SAT to SUBSET SUM. Imagine we are given an  $n$ -variable formula  $\phi$  in CNF with clauses  $C_1, \dots, C_m$  where each clause uses at most  $\ell$  variables. We create some  $(n + m)$ -digit numbers (in base  $\ell + 1$ ) that sum to a particular target  $t$  if and only if  $\phi$  is satisfiable. These numbers have one column corresponding to each variable, and one to each clause.

- For each variable  $v$ , create two integers  $n_{vt}$  and  $n_{vf}$  which are both 1 in the column corresponding to  $v$ .  $n_{vt}$  also has a 1 digit in the columns corresponding to clauses that are satisfied by  $v$  being true, and 0 digits elsewhere.  $n_{vf}$  also has a 1 digit in the columns corresponding to clauses that are satisfied by  $v$  being false, and 0 digits elsewhere.
- For each clause  $C_i$  containing  $\ell_i$  variables, create  $\ell_i - 1$  numbers that are 1 in the column corresponding to  $C_i$ , and 0 elsewhere.
- The target  $t$  has 1 digits in all the columns corresponding to variables, and a digit equal to  $\ell_i$  in the column corresponding to  $C_i$ .

Now, if there is a subsequence of integers summing to  $t$ , this must use exactly one of each of the pairs  $\{n_{vt}, n_{vf}\}$  (corresponding to each variable being either true or false) to make the first  $n$  columns correct. Also, each clause of  $\phi$  must have been satisfied by at least one variable, or the last  $m$  columns cannot be correct. Therefore, there is a satisfying assignment to  $\phi$  if and only if there is a subsequence of the numbers that sums to  $t$ .  $\square$

We will finally give two examples of NP-hardness proofs for problems of a graph-theoretic nature.

**Theorem 6.4.** CLIQUE is NP-complete.

*Proof.* CLIQUE is clearly in NP, as given a claimed clique on  $m$  vertices, it can be verified in time  $O(m^2)$ . To show that CLIQUE is NP-complete, we reduce 3-SAT to CLIQUE as follows. Given a 3-CNF formula with  $m$  clauses, we create an undirected graph  $G$  with at most  $3m$  vertices. We associate a set of at most three vertices with each clause  $C$ , labelled by  $C$  and the literals in  $C$ . We then connect every pair of vertices in  $G$ , with the exception of vertices in the same clause and vertices labelled with a variable and its negation (i.e.  $(x_i, \neg x_i)$  for some  $i$ ). See Figure 10 for an illustration. We now show that  $G$  has an  $m$ -clique if and only if  $\phi$  is satisfiable.

<sup>1</sup>The version of the problem we describe here would technically be better known as SUBSEQUENCE SUM.




Figure 10: Reducing satisfiability of the formula  $(x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee x_4) \wedge (x_2 \vee \neg x_3 \vee \neg x_4)$  to finding a 3-clique (one such clique is highlighted).

- (Satisfiable  $\Rightarrow$   $m$ -clique): Let  $w$  be a satisfying assignment to  $\phi$ . We form a set  $S$  of variables by picking a variable  $x_i$  from each clause  $C$  such that  $x_i = w_i$  and  $x_i$  satisfies  $C$  (this is always possible as  $w$  satisfies every clause of  $\phi$ ). Then there is an edge from every element of  $S$  to every other element, because each  $x_i$  is consistent with the fixed variable  $w_i$ , so we never have an edge  $(x_i, \neg x_i)$ .
- ( $m$ -clique  $\Rightarrow$  satisfiable): Any  $m$ -clique must include exactly one vertex from each clause, as there are no edges between vertices in the same clause. Also, the labels for the vertices must be consistent, as there are no edges between inconsistent labellings. Thus taking the labelling of each vertex in the clique (and assigning any other variables arbitrarily) gives a satisfying assignment for  $\phi$ .

As  $G$  can clearly be constructed from  $\phi$  in polynomial time, this completes the proof.  $\square$

The strategy used in the above proof of associating a subgraph with each fragment of the original formula, then connecting the subgraphs appropriately, is known as a *gadget construction* and is frequently a useful idea.

**Theorem 6.5.** HAMILTONIAN PATH is NP-complete.

*Proof.* HAMILTONIAN PATH is in NP because a claimed path visiting every vertex once can be checked in polynomial time. To show that the problem is NP-complete, we reduce SAT to HAMILTONIAN PATH. Given a formula  $\phi$  with  $n$  variables and  $m$  clauses, we will produce a (directed) graph  $G$  in polynomial time such that  $G$  has a Hamiltonian path if and only if  $\phi$  is satisfiable.

$G$  will have  $m$  nodes  $v_{C_i}$  corresponding to the clauses  $C_1, \dots, C_m$  of  $\phi$ ;  $n$  “chains” of  $2m + 1$  nodes each corresponding to the variables  $x_1, \dots, x_n$ ; and finally two special nodes  $v_{\text{in}}, v_{\text{out}}$ . A chain is a sequence of nodes  $v_i$  such that each pair  $v_i, v_{i+1}$  are connected by arcs in both directions.

We also have arcs from  $v_{\text{in}}$  to the start and end of the first chain, and from the start and end of the last chain to  $v_{\text{out}}$ ; and from the start and end of the  $j$ 'th chain to the start and end of the  $j + 1$ 'th chain, for each  $j$ . Finally, we have arcs corresponding to the clauses of  $\phi$ , as follows. If clause  $C$  contains the (unnegated) variable  $x_i$ , then we take two neighbouring nodes  $v_j, v_{j+1}$  in the  $i$ 'th chain, and put an arc from  $v_j$  to  $v_C$ , and from  $v_C$  to  $v_{j+1}$ . If  $C$  contains the negation  $\neg x_i$ , we instead put an arc from  $v_C$  to  $v_j$ , and from  $v_{j+1}$  to  $v_C$ . We never use the same node or arc for two different clauses (as each chain contains  $2m + 1$  nodes, we can always achieve this). This whole construction is illustrated in Figure 6.




Figure 11: Parts of the reduction from determining satisfiability of  $(x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee x_2 \vee \neg x_4)$  to finding a Hamiltonian path. The two pictures illustrate the chains without the clause nodes, and one clause node.

We now need to show that  $G$  has a Hamiltonian path if and only if  $\phi$  is satisfiable.

- (Satisfiable  $\Rightarrow$  Hamiltonian path): Assume  $\phi$  has satisfying assignment  $x$ . The path will start at  $v_{\text{in}}$  and traverse each chain in turn, finishing at  $v_{\text{out}}$ . The path traverses the  $i$ 'th chain from left to right if  $x_i = 1$ , and from right to left if  $x_i = 0$ . This is not quite a Hamiltonian path as it does not visit the “clause” vertices. But, as  $x$  is a satisfying assignment to  $\phi$ , for each clause  $C$  there must exist at least one chain such that the path can visit  $C$  en route from the start to end of that chain (i.e. the path is going in the right direction along that chain).
- (Hamiltonian path  $\Rightarrow$  satisfiable): First observe that any Hamiltonian path  $P$  must begin in  $v_{\text{in}}$  and end in  $v_{\text{out}}$ . We now claim that  $P$  needs to traverse each chain in order (each either from left to right, or right to left). If there were no clause nodes, this would be obvious; however, it might be the case that  $P$  goes from node  $v_i$  on the  $j$ 'th chain to a clause node  $C$ , then goes from  $C$  to a different chain  $k$ . Nevertheless, this can never give a Hamiltonian path: when  $P$  reaches the neighbouring node  $v_{i+1}$  (or  $v_{i-1}$ ) it will get stuck, as there are no arcs leaving this neighbour to nodes which have not already been visited. Hence forming a bit-string  $x$  by taking  $x_i = 1$  if the  $i$ 'th chain is traversed from left to right, and  $x_i = 0$  if the  $i$ 'th chain is traversed from right to left, gives a satisfying assignment to  $\phi$ .

□

It may be already apparent that, although NP-completeness proofs can usually be readily checked, devising a suitable reduction to prove NP-completeness can require some ingenuity. This is perhaps appropriate given the definition of NP.

## 6.1 Historical notes and further reading

The theory of NP-completeness took off in 1972 when Karp proved that the decision versions of many practically important optimisation problems are NP-complete; literally thousands of other problems have since been proven NP-complete, and this terminology has now entered the standard

scientific vocabulary. Garey and Johnson is the standard reference for NP-completeness results (albeit now somewhat outdated).

Could it be the case that *all* problems in NP are either in P or NP-complete? Assuming that  $P \neq NP$ , it has been shown that this cannot be true: there must exist problems in NP which are neither in P nor NP-complete (“Ladner’s theorem”, see Papadimitriou chapter 14). However, very few natural candidate problems are known which are thought to belong to  $NP \setminus P$ . One example is INTEGER FACTORISATION. Another is GRAPH ISOMORPHISM: given the description of two undirected graphs  $G$  and  $H$ , is there a permutation  $\pi$  of the vertices of  $G$  such that  $\pi(G) = H$ ?



## 7 The P vs. NP problem and beyond

We know that  $P \subseteq NP$ . The million-dollar question<sup>1</sup> is whether this inclusion is strict. This question has vast practical significance (many problems which people want to solve in practice turn out to be NP-complete, some of which are given in Section 6), but it also has significant theoretical importance, as it essentially asks whether mathematics can be axiomatised in the following concrete sense.

Recall from Section 3.8 that the problem of determining whether a statement about the natural numbers can be proven from the axioms of Peano arithmetic is undecidable. In practice, we may be satisfied with the following weaker notion of axiomatisation. Given a statement about the natural numbers, does it have a proof (i.e. a derivation of the statement from the axioms) of a reasonable length<sup>2</sup> in Peano arithmetic? One can formalise this as the following decision problem, which we call THEOREM. Given a statement  $S$  of length  $n$  about the natural numbers, and an integer  $k$ , does  $S$  have a proof from the axioms of Peano arithmetic of length at most  $n^k$ ?

**Theorem 7.1.** THEOREM is NP-complete.

*Proof.* It is clear that THEOREM is in NP: given a claimed proof of  $S$  in any “sensible” formal system, we can check it in time polynomial in the length of the proof. To see that it is NP-hard, we reduce 3-SAT to THEOREM using a process known as arithmetisation. Given a 3-CNF formula  $\phi$ , each clause  $C$  containing variables  $(x_i, x_j, x_k)$  is replaced by the product  $(1 - z_i z_j z_k)$ , where  $z_i = x_i$  if  $x_i$  appears negated in  $C$ , and  $z_i = (1 - x_i)$  otherwise. We then form a polynomial  $P_\phi$  by taking the product of these polynomials over all clauses  $C$ . For example,

$$(x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee x_3) \mapsto (1 - (1 - x_1)x_2(1 - x_3))(1 - x_1(1 - x_3)).$$

Now  $\phi$  will have a satisfying assignment  $x$  if and only if there exist  $x_1, \dots, x_n \in \{0, 1\}$  such that  $P_\phi(x_1, \dots, x_n) = 1$ , which is a simple statement about the natural numbers.  $\square$

Thus, if  $P = NP$ , we can determine whether statements have short proofs. More is true: if  $P = NP$ , we can even *find* such short proofs.

**Theorem 7.2.** If  $P = NP$ , for every language  $\mathcal{L} \in NP$  there exists a polynomial-time Turing machine  $M$  that, on input  $x \in \mathcal{L}$ , outputs a certificate for  $x$  with respect to  $\mathcal{L}$  and  $M$ . That is, there exists a polynomial-time Turing machine  $V_{\mathcal{L}, M}$  such that  $V_{\mathcal{L}, M}(M(x)) = 1$  if and only if  $x \in \mathcal{L}$ .

*Proof.* We first show that the above claim holds for SAT, by showing that given access to a polynomial-time algorithm  $\mathcal{A}$  for deciding whether some input formula  $\phi$  is satisfiable, we can find a satisfying assignment  $w$  (if one exists) in polynomial time. Our algorithm first runs  $\mathcal{A}$  to determine satisfiability of  $\phi$ . If it is satisfiable, we consider each of the two formulae  $\phi_0, \phi_1$  produced by fixing  $x_1 = 0, x_1 = 1$ , and use  $\mathcal{A}$  to test their satisfiability. At least one of these formulae will be satisfiable; whichever one of them is, we fix  $x_1$  appropriately and continue with  $x_2$ , etc. We eventually fix all the variables and have determined a satisfying assignment to  $\phi$ .

In order to see that the above strategy works for *any*  $\mathcal{L} \in NP$ , observe that the proof of the Cook-Levin theorem maps a certificate for  $\mathcal{L}$  (i.e. a computational path of an NDTM) to a satisfying assignment  $w$  of a boolean formula in such a way that the certificate can be retrieved from  $w$ .  $\square$

---

<sup>1</sup>Literally: the P vs. NP problem is one of the Clay Mathematics Institute’s Millennium Prize problems, and solving it carries a prize of \$1m.

<sup>2</sup>After all, can a claim really be said to be “true” if the proof is too long to check?

Thus, if  $P = NP$ , there would exist a polynomial-time algorithm for finding short proofs of theorems: informally speaking, mathematical creativity would be automated. It would seem somewhat surprising if this were possible, and hence most effort to date has concentrated on proving the conjecture that  $P \neq NP$ . Unfortunately, there are several substantial barriers known to proving this conjectured separation of complexity classes.

## 7.1 The relativisation barrier

We previously used the technique of diagonalisation to prove both the undecidability of the halting problem and the time hierarchy theorem. It is natural to hope that this technique could also be useful to prove  $P \neq NP$ . However, we will now see a significant problem with using this technique to attack the  $P$  vs.  $NP$  question, which is based on the concept of oracles.

An oracle Turing machine is a standard Turing machine  $M$  which is given access to an *oracle* which (via some magical process!) solves the decision problem for some language  $\mathcal{L}$ , where  $\mathcal{L}$  may not be solvable by  $M$  without the oracle. Formally, an oracle Turing machine  $M$  has three additional states {QUERY, YES, NO}, and an additional read/write *oracle tape*. The operation of  $M$  depends on the input  $x$  and an additional language  $\mathcal{L}$  which is called the oracle for  $M$ . If  $M$  enters the QUERY state during the computation, in the next step  $M$  enters the state YES if the contents  $q$  of the oracle tape satisfy  $q \in \mathcal{L}$ , and enters the state NO if  $q \notin \mathcal{L}$ . Let  $M^{\mathcal{L}}(x)$  denote the output of  $M$  on input  $x$  with an oracle for  $\mathcal{L}$ . Note that  $M$  can be either a deterministic or nondeterministic Turing machine.

We also define  $P^{\mathcal{L}}$  (resp.  $NP^{\mathcal{L}}$ ) to be the set of languages which can be decided by a polynomial-time deterministic (resp. nondeterministic) Turing machine with oracle access to  $\mathcal{L}$ .

Some simple examples to illustrate the concept of oracles are:

- For any  $\mathcal{L} \in P$ ,  $P^{\mathcal{L}} = P$ . The inclusion  $P \subseteq P^{\mathcal{L}}$  holds because allowing oracle calls can only increase the set of languages that can be decided, while the inclusion  $P^{\mathcal{L}} \subseteq P$  holds because a polynomial-time Turing machine can replace each call to an oracle for  $\mathcal{L}$  by simply deciding  $\mathcal{L}$  directly.
- Let EC be the language  $\{(\underline{M}, x, 1^n) : M \text{ outputs } 1 \text{ on } x \text{ within } 2^n \text{ steps}\}$ . We show that  $P^{EC} = NP^{EC} = EXP$ . On the one hand, calling the oracle for EC once on an input of size  $n^k$  allows an arbitrary exponential-time computation to be performed, so  $EXP \subseteq P^{EC}$ . On the other, given exponential time, a deterministic Turing machine can simulate each of the possible calls made to the EC oracle by an NDTM running in polynomial time (there can be exponentially many possible calls, each of which can be simulated in exponential time), so  $NP^{EC} \subseteq EXP$ . Hence  $EXP \subseteq P^{EC} \subseteq NP^{EC} \subseteq EXP$ . We say that “relative to EC,  $P = NP$ ”.

We now show that, relative to oracles, the  $P$  vs.  $NP$  question has two contradictory resolutions.

**Theorem 7.3** (Baker, Gill and Solovay). *There exist languages  $\mathcal{A}$ ,  $\mathcal{B}$  such that  $P^{\mathcal{A}} = NP^{\mathcal{A}}$ , but  $P^{\mathcal{B}} \neq NP^{\mathcal{B}}$ .*

*Proof.* We have already proven the first part by taking  $\mathcal{A}$  to be the previously discussed language EC. The second part is more technical. For any language  $\mathcal{B} \subseteq \{0, 1\}^*$ , let  $U_{\mathcal{B}}$  be the language

$$U_{\mathcal{B}} = \{1^n : \mathcal{B} \text{ contains a string of length } n\}.$$

For every  $\mathcal{B}$ ,  $U_{\mathcal{B}}$  is clearly in  $\text{NP}^{\mathcal{B}}$ , as the NP machine can guess a string  $x \in \mathcal{B}$  of length  $n$ , which can be checked using the oracle for  $\mathcal{B}$ . We now show that there exists a language  $\mathcal{B}$  such that  $U_{\mathcal{B}} \notin \text{P}^{\mathcal{B}}$ , so  $\text{P}^{\mathcal{B}} \neq \text{NP}^{\mathcal{B}}$ .

We construct  $\mathcal{B}$  via a choosing process using infinitely many stages; each stage will determine whether a finite-sized set of strings is in  $\mathcal{B}$ . Initially,  $\mathcal{B}$  is empty. For each stage  $i$ , let  $M_i$  be the oracle deterministic Turing machine represented by (the binary expansion of)  $i$ . For technical reasons, we will need to assume that each Turing machine  $M$  is represented by infinitely many such strings  $i$ ; this can easily be achieved by, for example, allowing representations to be padded by arbitrarily many trailing zeroes.

For each  $i$ , choose  $n$  to be larger than the length of any string in  $\mathcal{B}$ . We will ensure that  $M_i^{\mathcal{B}}$  does not decide  $U_{\mathcal{B}}$  in time at most  $2^n/10$ . To do so, we run  $M_i$  on input  $1^n$  for  $2^n/10$  steps. Whenever  $M_i$  makes oracle queries corresponding to a string which the choosing process has determined to be in  $\mathcal{B}$  (or not), the oracle answers consistent with this. When  $M_i$  queries a string  $q$  for which it has not yet been chosen whether  $q \in \mathcal{B}$ , the oracle answers that  $q \notin \mathcal{B}$ , and fixes  $q \notin \mathcal{B}$ . After  $M_i$  has finished its computation, it must either accept or reject  $1^n$ ; we will ensure that in either case its answer is incorrect. If  $M_i$  accepts  $1^n$ , we declare that all remaining strings of length  $n$  are not in  $\mathcal{B}$ , so  $1^n \notin U_{\mathcal{B}}$ . On the other hand, if  $M_i$  rejects  $1^n$ , we find a string  $x$  of length  $n$  not queried by  $M_i$  (which can be done because  $M_i$  only queried at most  $2^n/10$  strings) and declare  $x \in \mathcal{B}$ , so  $1^n \in U_{\mathcal{B}}$ . As this holds for all  $i$ ,  $U_{\mathcal{B}}$  is not decided by any deterministic oracle Turing machine running in time  $2^n/10$ . As every polynomial is smaller than  $2^n/10$  for large enough  $n$ , and every Turing machine  $M$  is represented by infinitely many  $i$ , if  $M$  runs in polynomial time it cannot decide  $U_{\mathcal{B}}$ . Thus  $U_{\mathcal{B}} \notin \text{P}^{\mathcal{B}}$ .  $\square$

The significance of Theorem 7.3 is that it implies that proof techniques which are “blind” to the existence of oracles will not be able to resolve the P vs. NP problem. Indeed, the proofs of the undecidability of the Halting Problem and the Time Hierarchy Theorem which we saw previously rely only on the following two properties:

- The countability of Turing machines (i.e. the existence of an effective representation of these machines by binary strings);
- The ability of one Turing machine to simulate another efficiently (i.e. without significant time overhead).

Oracle Turing machines also satisfy these properties, but we have seen that the P vs. NP question cannot be resolved for such machines. Any resolution of the problem must therefore use some additional property of standard Turing machines.

Later on, following the introduction of concepts concerned with randomness and circuit complexity, we discuss a different roadblock to proving  $\text{P} \neq \text{NP}$ , called the natural proofs barrier.

## 7.2 Complementation of languages and co-NP

For any language  $\mathcal{L} \subseteq \Sigma^*$  we can define the complement of  $\mathcal{L}$ ,  $\bar{\mathcal{L}} = \{x \in \Sigma^*, x \notin \mathcal{L}\}$ . In the case of a decision problem, we define the notion of complementation as follows: given a decision problem  $D$ , the complement of  $D$  is the problem whose answer is “yes” whenever  $D$ ’s answer is “no”, and vice versa. For example, the complement of the 3-SAT problem is: given a 3-CNF formula  $\phi$ , is it *not* satisfiable – i.e. is it the case that there is no  $x$  such that  $\phi(x) = 1$ ? Note that the language corresponding to the complement of 3-SAT is not quite the same as the complement of the language

corresponding to 3-SAT. The former is the set of unsatisfiable 3-CNF formulae, while the latter is the set of bit-strings which do not encode satisfiable 3-CNF formulae. Usually this distinction will not be important.

For any complexity class  $C$ , we can thus define the complexity class  $\text{co-}C$  as the set  $\{\bar{\mathcal{L}} : \mathcal{L} \in C\}$ . It is usually not the case that  $\text{co-}C$  is the complement of  $C$ . Indeed, it is immediate that, for example,  $\text{co-P} = \text{P}$ : a polynomial-time algorithm  $\mathcal{A}$  for deciding  $\mathcal{L}$  can be modified to give a polynomial-time algorithm  $\bar{\mathcal{A}}$  for deciding  $\bar{\mathcal{L}}$  by inverting the output of  $\mathcal{A}$ . However, in the case of nondeterministic classes the question is not so easy: in particular, it is not known whether  $\text{co-NP} = \text{NP}$ . Intuition behind this can be gained from the 3-SAT problem discussed above. Given a satisfiable formula  $\phi$ , a satisfying assignment  $x$  gives a certificate that it is indeed satisfiable, which can be checked efficiently. On the other hand, if  $\phi$  is not satisfiable by any assignment  $x$ , it is not clear how to exhibit a certificate of this which can be checked efficiently.

There is an alternative “operational” definition of  $\text{co-NP}$  which corresponds more closely (and perhaps intuitively) to our original definition of  $\text{NP}$ .

**Definition 7.1.** For every  $\mathcal{L} \subseteq \{0, 1\}^*$ , we say that  $\mathcal{L} \in \text{co-NP}$  if there exists a polynomial  $p : \mathbb{N} \rightarrow \mathbb{N}$  and a polynomial-time Turing machine  $M$  such that, for all  $x \in \{0, 1\}^*$ ,  $x \in \mathcal{L}$  if and only if for all  $u \in \{0, 1\}^{p(|x|)}$ ,  $M(x, u) = 1$ .

Note the change from “there exists” to “for all”. We now check that these two definitions are indeed equivalent. If  $\mathcal{L}$  satisfies Definition 7.1, then  $x \in \bar{\mathcal{L}}$  if and only if there exists a  $u$  such that  $M(x, u) = 0$ ; hence  $\bar{\mathcal{L}} \in \text{NP}$  (the  $\text{NP}$  verifier simply runs  $M$  and inverts the result). But if  $\bar{\mathcal{L}} \in \text{NP}$ , there exists  $M'$  such that  $x \in \mathcal{L} \Leftrightarrow M'(x, u) = 0$  for all  $u$ , hence  $\mathcal{L} \in \text{co-NP}$  by running  $M'$  and inverting the result.

Just as with  $\text{NP}$ ,  $\text{co-NP}$  has complete problems, i.e. languages  $\mathcal{L} \in \text{co-NP}$  such that for all  $\mathcal{L}' \in \text{co-NP}$ ,  $\mathcal{L}' \leq_P \mathcal{L}$ . A natural example is the problem  $\text{TAUTOLOGY}$ : given a boolean formula  $\phi$ , expressed in  $\text{DNF}$ , is  $\phi$  satisfied by *every* assignment to its variables? Clearly,  $\phi$  is a tautology if and only if  $\neg\phi$  is not satisfiable.

**Lemma 7.4.**  $\text{TAUTOLOGY}$  is  $\text{co-NP}$ -complete.

*Proof.* We need to show that, for  $\mathcal{L} \in \text{co-NP}$ ,  $\mathcal{L} \leq_P \text{TAUTOLOGY}$ . If  $\mathcal{L} \in \text{co-NP}$ , there is a polynomial-time reduction  $R$  from  $\bar{\mathcal{L}}$  to  $\text{SAT}$ , so for all  $x \in \{0, 1\}^*$ ,  $x \in \bar{\mathcal{L}}$  if and only if  $R(x)$  is satisfiable. Hence  $x \in \mathcal{L}$  if and only if  $\neg R(x)$  is satisfied by every assignment to its variables. Observe that, using De Morgan’s law,  $\neg R(x)$  can be written in  $\text{DNF}$  as required<sup>1</sup> by our definition of  $\text{TAUTOLOGY}$ .  $\square$

We have the following characterisation of all  $\text{co-NP}$ -complete problems.

**Theorem 7.5.**  $\mathcal{L}$  is  $\text{co-NP}$ -complete if and only if  $\bar{\mathcal{L}}$  is  $\text{NP}$ -complete.

*Proof.* We will show that, if  $\bar{\mathcal{L}}$  is  $\text{NP}$ -complete,  $\mathcal{L}$  is  $\text{co-NP}$ -complete (the other direction is similar). If  $\bar{\mathcal{L}} \in \text{NP}$ ,  $\mathcal{L} \in \text{co-NP}$  by definition. We also need to show that if  $\bar{\mathcal{L}}$  is  $\text{NP}$ -hard, then  $\mathcal{L}$  is  $\text{co-NP}$ -hard. Consider an arbitrary language  $\mathcal{A} \in \text{co-NP}$ . As  $\bar{\mathcal{L}}$  is  $\text{NP}$ -hard, and  $\bar{\mathcal{A}} \in \text{NP}$ , there exists a polynomial-time computable function  $R : \{0, 1\}^* \rightarrow \{0, 1\}^*$  such that  $x \in \bar{\mathcal{A}} \Leftrightarrow R(x) \in \bar{\mathcal{L}}$ . Thus  $x \in \mathcal{A} \Leftrightarrow R(x) \in \mathcal{L}$ , so  $\mathcal{L}$  is  $\text{co-NP}$ -hard.  $\square$

<sup>1</sup>The alert reader may have spotted that, just as  $\text{SAT}$  remains  $\text{NP}$ -complete if we relax the restriction that the input formula is in  $\text{CNF}$ ,  $\text{TAUTOLOGY}$  remains  $\text{co-NP}$ -complete if we allow input formulae which are not in  $\text{DNF}$ .

Could it be possible that  $\text{NP} = \text{co-NP}$ ? At the moment, we cannot rule this out, though such a result would seem almost as remarkable as  $\text{P} = \text{NP}$ : it would (informally) imply that any problem which has a short certificate of having a positive solution must also have a short certificate of having a negative solution. It is also unknown whether or not  $\text{P} = \text{NP} \cap \text{co-NP}$ ; while this would also be surprising, there are very few examples of problems known to be in  $\text{NP} \cap \text{co-NP}$  but not known to be in  $\text{P}$ .

### 7.3 The polynomial hierarchy

We have seen that the classes  $\text{NP}$  and  $\text{co-NP}$  have natural definitions in terms of “there exists” and “for all” quantifiers, respectively:

- $\mathcal{L} \in \text{NP}$ :  $x \in \mathcal{L}$  if and only if  $\exists w \in \{0, 1\}^{\text{poly}(|x|)}, M(x, w) = 1$ .
- $\mathcal{L} \in \text{co-NP}$ :  $x \in \mathcal{L}$  if and only if  $\forall w \in \{0, 1\}^{\text{poly}(|x|)}, M(x, w) = 1$ .

If we combine these two types of quantifier, we get a potentially more general set of complexity classes known as the polynomial hierarchy<sup>1</sup>.

**Definition 7.2.** For every  $\mathcal{L} \subseteq \{0, 1\}^*$ , and for integer  $i \geq 1$ , we say that  $\mathcal{L} \in \Sigma_i$  if there exists a polynomial  $p : \mathbb{N} \rightarrow \mathbb{N}$  and a polynomial-time Turing machine  $M$  such that, for all  $x \in \{0, 1\}^*$ ,

$$x \in \mathcal{L} \Leftrightarrow \exists u_1 \in \{0, 1\}^{p(|x|)} \forall u_2 \in \{0, 1\}^{p(|x|)} \dots Q_i u_i \in \{0, 1\}^{p(|x|)}, M(x, u_1, \dots, u_i) = 1,$$

where the quantifier  $Q_i$  is  $\exists$  (if  $i$  is odd) or  $\forall$  (if  $i$  is even). The polynomial hierarchy is the set  $\text{PH} = \bigcup_{i \geq 1} \Sigma_i$ .

To help digest this definition, consider the following problem, called **LARGEST CLIQUE**. Given a graph  $G$  and an integer  $k$ , is the largest clique within  $G$  of size exactly  $k$ ? This problem is in the class  $\Sigma_2$ . To see this, observe that a graph has a largest clique of size  $k$  if and only if: there exists a set  $S$  of vertices such that  $S$  is of size  $k$ , and  $S$  is a clique; and for all sets  $S'$  of size  $k + 1$ ,  $S'$  is not a clique.

Recall that we showed that the problem **CLIQUE** (given a graph  $G$  and an integer  $k$ , does  $G$  contain a clique on  $k$  vertices?) was  $\text{NP}$ -complete. However, if **LARGEST CLIQUE** were contained within  $\text{NP}$ , there would exist an efficiently checkable certificate that there are no cliques in  $G$  of size at least  $k + 1$ , and it is not clear what such a certificate would look like. Thus, similarly to the intuition that  $\text{NP} \neq \text{co-NP}$ , it is believed that  $\Sigma_i \neq \Sigma_{i+1}$  for all  $i$ .

Just as with  $\text{NP}$  and  $\text{co-NP}$ , each level of the polynomial hierarchy has a natural complete problem, a “quantified” variant of **SAT**. Let  $\phi$  be a boolean formula on  $n$  variables  $x_1, \dots, x_n$  partitioned into sets  $X_1, \dots, X_i$ . Then a  $\text{QSAT}_i$  instance is an expression

$$\exists X_1 \forall X_2 \exists X_3 \dots Q_i X_i \phi,$$

where  $Q_i$  is  $\exists$  (if  $i$  is odd) or  $\forall$  (if  $i$  is even). For example, the following expression is a  $\text{QSAT}_3$  instance:

$$\exists x_1, x_2 \forall x_3 \exists x_4, x_5 (x_1 \vee \neg x_2) \wedge (\neg x_3 \vee x_5) \wedge (\neg x_1 \vee x_4).$$

The  $\text{QSAT}_i$  problem is then simply to determine whether such an expression evaluates to true. Observe that  $\text{QSAT}_1$  is just the **SAT** problem. More generally, we have the following theorem, which we will not prove but should not be too surprising given the definitions of  $\text{QSAT}_i$  and  $\Sigma_i$ .

<sup>1</sup>Short for “polynomial-time hierarchy”.

**Theorem 7.6.**  $QSAT_i$  is  $\Sigma_i$ -complete.

This allows us to give the following alternative and concise definition of the polynomial hierarchy in terms of oracle machines.

**Theorem 7.7.**  $\Sigma_1 = \text{NP}$ , and for  $i \geq 2$ ,  $\Sigma_i = \text{NP}^{QSAT_{i-1}}$ .

*Proof idea.* To see the essential ideas, we will prove the special case that  $\Sigma_2 = \text{NP}^{\text{SAT}}$ .

- ( $\Sigma_2 \subseteq \text{NP}^{\text{SAT}}$ .) Let  $\mathcal{L} \in \Sigma_2$ . Then there exists a polynomial  $p : \mathbb{N} \rightarrow \mathbb{N}$  and a polynomial-time Turing machine  $M$  such that, for all  $x \in \{0, 1\}^*$ ,

$$x \in \mathcal{L} \Leftrightarrow \exists u_1 \in \{0, 1\}^{p(|x|)} \forall u_2 \in \{0, 1\}^{p(|x|)}, M(x, u_1, u_2) = 1.$$

Our nondeterministic algorithm with access to an oracle for SAT will proceed as follows. On input  $x$ , the algorithm guesses a string  $u_1$ , and would like to determine whether there exists  $u_2 \in \{0, 1\}^{p(|x|)}$  such that  $M(x, u_1, u_2) = 0$ . By the Cook-Levin theorem, this can be written as a SAT instance. The algorithm therefore uses its oracle for SAT, which outputs YES if there exists  $u_2 \in \{0, 1\}^{p(|x|)}$  such that  $M(x, u_1, u_2) = 0$ , and accepts if and only if the answer is NO. Thus it accepts if and only if  $\exists u_1 \in \{0, 1\}^{p(|x|)} \forall u_2 \in \{0, 1\}^{p(|x|)}, M(x, u_1, u_2) = 1$ .

- ( $\text{NP}^{\text{SAT}} \subseteq \Sigma_2$ .) Let  $\mathcal{L} \in \text{NP}^{\text{SAT}}$ , and let  $M$  be an NDTM which decides  $\mathcal{L}$  with access to an oracle for SAT. We first show that there is a machine  $N$  which also decides  $\mathcal{L}$ , makes one query to the oracle, and accepts if and only if the answer to the oracle query is NO. On input  $x$ ,  $N$  guesses all of the nondeterministic choices of  $M$ , all of the oracle questions, and all of the answers. For each question that was answered with YES,  $N$  guesses a satisfying assignment to check that the guess was correct. At the end,  $N$  still has a set of oracle questions (i.e. formulae)  $\phi_1, \dots, \phi_k$  to which the oracle would have replied with NO.  $N$  then asks the oracle whether the formula  $\phi_1 \vee \dots \vee \phi_k$  is satisfiable, and accepts if and only if the answer is NO.

We now observe that  $N$  accepts (i.e.  $x \in \mathcal{L}$ ) if and only if there exist a set of nondeterministic choices, oracle questions and answers, such that there is no satisfying assignment to  $\phi_1 \vee \dots \vee \phi_k$ . Thus  $\mathcal{L} \in \Sigma_2$  as required. □

## 7.4 Historical notes and further reading

Perhaps the first to recognise the connection between the P vs. NP problem and the foundations of mathematics was Gödel<sup>1</sup> in 1956. Baker, Gill and Solovay proved their result in 1975; a proof is given in Arora-Barak chapter 3, including a brief discussion of the conceptual connection to independence results in mathematical logic. Papadimitriou chapter 17 has an extensive discussion of the polynomial hierarchy, which was defined by Meyer and Stockmeyer in 1972; there are also some good lecture notes by Luca Trevisan<sup>2</sup>.

<sup>1</sup><http://rjlipton.wordpress.com/the-gdel-letter/>

<sup>2</sup><http://www.cs.berkeley.edu/~luca/cs278-08/lecture04.pdf>

## 8 Space complexity

So far we have studied computations which are restricted with respect to the time that they use. Another important resource to consider is space. In the Turing machine model, this corresponds to the number of cells used on the tape.

In particular, we would like to study computations which use less space than the size of the input. In order for this to make sense, we modify the Turing machine model as follows.

**Definition 8.1.** A space-bounded (deterministic/nondeterministic) Turing machine  $M$  is a multi-tape Turing machine with a special read-only input tape, on which the input  $x$  is placed prior to the start of the computation. All other tapes are known as work tapes. The output  $M(x)$  is placed on one of these work tapes, which is known as the output tape.

Let  $s : \mathbb{N} \rightarrow \mathbb{N}$  and  $\mathcal{L} \subseteq \{0,1\}^*$ . We say that  $\mathcal{L} \in \text{SPACE}(s(n))$  if there is a constant  $c > 0$  and a deterministic Turing machine  $M$  which decides  $\mathcal{L}$ , such that at most  $cs(n)$  locations on  $M$ 's work tapes are ever scanned by  $M$ 's head during the computation, on any input of size  $n$ . Similarly,  $\mathcal{L} \in \text{NSPACE}(s(n))$  if there is an NDTM  $M$  which decides  $\mathcal{L}$ , under the same constraints for any of  $M$ 's nondeterministic choices. We usually restrict to *space-constructible* functions  $s(n)$ , i.e. functions  $s(n)$  for which there exists a Turing machine that computes  $s(|x|)$  in space  $O(s(|x|))$  on input  $x$ . As with time-constructible functions, this is a very minor restriction as all “sensible” functions (e.g.  $\lceil \log n \rceil$ ,  $n$ ,  $2^n$ ) are space-constructible.

By analogy with their time-complexity equivalents, we define  $\text{PSPACE} = \bigcup_{c>0} \text{SPACE}(n^c)$ ,  $\text{NSPACE} = \bigcup_{c>0} \text{NSPACE}(n^c)$ . Two classes which do not have a sensible time-complexity analogue are

$$\text{L} = \text{SPACE}(\log n), \text{NL} = \text{NSPACE}(\log n).$$

Observe that in these cases the space used is significantly less than the size of the input; the corresponding time complexity classes would correspond to algorithms which do not have time to read the whole input, which are usually uninteresting<sup>1</sup>. We can thus imagine an L computation as one where the size of the data on which the computer is operating is much larger than the size of the computer itself. An intuitive example of this phenomenon from daily experience is the use of the Internet: the Internet is too large to be stored within our computers, but we can access it and retrieve data, on which we can perform computations.

As with the case of NP, the class NL has an alternative definition in terms of certificates, which we formalise as follows.

**Definition 8.2.** NL is the class of languages  $\mathcal{L} \subseteq \{0,1\}^*$  such that there exists a polynomial  $p : \mathbb{N} \rightarrow \mathbb{N}$  and a space-bounded Turing machine  $M$  such that:

- $M$  has an additional certificate tape which contains a witness  $w$ . The certificate tape is read-only and scanned once, left to right;
- for all  $x \in \{0,1\}^*$ ,  $x \in \mathcal{L}$  if and only if there exists  $w \in \{0,1\}^{p(|x|)}$  such that  $M(x, w) = 1$ ;
- $M$  uses at most  $O(\log |x|)$  space on its work tapes for all  $x$ .

It should be clear that this alternative definition is equivalent, as making a sequence of nondeterministic choices can be viewed as reading the bits of a certificate one by one, in order. The

---

<sup>1</sup>However, at the end of the course we will see a model in which this idea does make sense.

“read-once” restriction is necessary: if it is removed, the class of languages decided by such machines becomes equal to NP.

We observe that for any  $f : \mathbb{N} \rightarrow \mathbb{N}$ ,  $\text{DTIME}(f(n)) \subseteq \text{SPACE}(f(n))$ , as no machine using  $T$  computational steps can visit more than  $T$  tape cells. However, we can say somewhat more.

**Theorem 8.1.** *For any function  $f : \mathbb{N} \rightarrow \mathbb{N}$  such that  $f(n) \geq \log_2 n$ ,*

$$\text{DTIME}(f(n)) \subseteq \text{NTIME}(f(n)) \subseteq \text{SPACE}(f(n)) \subseteq \text{NSPACE}(f(n)) \subseteq \text{DTIME}(2^{O(f(n))}).$$

*Proof.* The first and third inclusions are obvious. The second is based on the exponential-time simulation of NDTMs by DTMs given in Theorem 5.2. Given an NDTM with two transition functions that runs in time  $T$ , we simulate it by a DTM which tries each of the possible  $2^T$  sequences of choices of transition functions in turn. To iterate through each of these transition functions requires only  $O(T)$  space (to store a counter).

To prove the final inclusion, we use the concept of a *configuration graph*. For every space  $s(n)$  deterministic or nondeterministic Turing machine  $M$  and input  $x \in \{0, 1\}^*$  such that  $|x| = n$ , the configuration graph  $G_{M,x}$  of  $M$  on input  $x$  is the directed graph whose nodes are all possible configurations  $c$  of  $M$  such that the input is  $x$  and the work tapes have at most  $s(n)$  non-blank cells in total.  $G_{M,x}$  has an arc  $c \rightarrow c'$  if  $M$ 's transition function can go from  $c$  to  $c'$ . If  $M$  is deterministic,  $G_{M,x}$  has out-degree 1 (i.e. is a path), and if  $M$  is nondeterministic,  $G_{M,x}$  has out-degree at most 2. Let the starting configuration be  $c_S$ , and assume that there is a unique configuration  $c_A$  on which  $M$  accepts (this is without loss of generality as we can modify  $M$  to erase the contents of all its work tapes before accepting). Then  $M$  accepts  $x$  if and only if there is a path from  $c_S$  to  $c_A$ .

Observe that, for some constant  $c$  (depending on  $M$ )  $G_{M,x}$  has at most  $n2^{cs(n)}$  nodes, as a configuration is completely defined by  $M$ 's state, the head positions<sup>1</sup> and the contents of the work tapes. Therefore, by enumerating all possible configurations, the graph  $G_{M,x}$  can be constructed in time  $2^{O(s(n))}$ . We can then test whether there is a path from  $c_S$  to  $c_A$  using the algorithm for PATH, which runs in time  $\text{poly}(2^{O(s(n))}) = 2^{O(s(n))}$ .  $\square$

We also have the following analogue of the Time Hierarchy Theorem for space complexity.

**Theorem 8.2** (Space Hierarchy Theorem). *If  $f$  and  $g$  are space-constructible functions such that  $f(n) = o(g(n))$ , then  $\text{SPACE}(f(n)) \subset \text{SPACE}(g(n))$ .*

The proof is the same as for the Time Hierarchy Theorem, but note that the result is somewhat tighter. The reason is that it is possible to design a universal Turing machine with only a constant factor space overhead.

## 8.1 Polynomial space

Intuitively, space appears to be a more powerful resource than time, as space can be reused whereas time cannot. Remarkably, it is nevertheless an open question whether  $\text{P} = \text{PSPACE}$ . However, as (by Theorem 8.1), we know that  $\text{NP} \subseteq \text{PSPACE}$ , it seems unlikely that this equality holds. In fact, it turns out that we can solve problems even more general than SAT in PSPACE.

A quantified boolean formula is a formula of the form

$$Q_1 x_1 Q_2 x_2 \dots Q_n x_n \phi(x_1, \dots, x_n),$$

---

<sup>1</sup>The “ $n$ ” comes from the input tape head position. As  $s(n) \geq \log_2 n$ , this can be absorbed into  $2^{O(s(n))}$ .



where each  $Q_i \in \{\forall, \exists\}$ , and  $\phi$  is a standard boolean formula. The TQBF problem (“true/totally quantified boolean formula”) is defined as follows: given a quantified boolean formula  $\psi$ , determine whether  $\psi$  is true. Observe that this is a generalisation of the QSAT<sub>*i*</sub> problem which we previously encountered: rather than the number of quantifiers being at most *i*, for some fixed *i*, now there can be up to *n* quantifiers. The above formula has all its quantifiers at the start (so-called “prenex normal form”). One could conceive of more general quantified boolean formulae with quantifiers appearing throughout the formula, e.g.

$$\exists x_1 (x_1 \vee \forall x_2, x_3 (\neg x_2 \wedge x_3)).$$

In fact, this is not a generalisation as all such formulae can be converted into prenex normal form in polynomial time. The quantifiers  $Q_1, \dots, Q_n$  may be assumed to alternate between  $\forall$  and  $\exists$  by introducing dummy variables. That is,

$$\exists x_1 \exists x_2 \phi(x_1, x_2) = \exists x_1 \forall y_1 \exists x_2 \phi(x_1, x_2),$$

etc. Observe that SAT is the special case of TQBF where each quantifier  $Q_i = \exists$ . That is, a boolean formula  $\phi$  is satisfiable if and only if

$$\exists x_1 \exists x_2 \dots \exists x_n \phi(x_1, \dots, x_n).$$

**Theorem 8.3.** *TQBF is PSPACE-complete.*

In the statement of this theorem, PSPACE-completeness is used exactly in the sense that NP-completeness was before. Explicitly, we say that a language  $\mathcal{A}$  is PSPACE-complete if  $\mathcal{A} \in \text{PSPACE}$ , and for all  $\mathcal{B} \in \text{PSPACE}$ ,  $\mathcal{B} \leq_P \mathcal{A}$ .

*Proof.* We first show that TQBF is in PSPACE. Let  $\psi = Q_1 x_1 Q_2 x_2 \dots Q_n x_n \phi(x_1, \dots, x_n)$  be a quantified boolean formula, and let the size of  $\phi$  be *m*. We give a recursive algorithm  $\mathcal{A}$  which can decide whether  $\psi$  is true in space  $\text{poly}(n, m)$ .

For  $b \in \{0, 1\}$ , let  $\psi_{x_1=b}$  denote the quantified boolean formula obtained by removing the first quantifier  $Q_1$  and replacing  $x_1$  with *b* throughout the formula  $\phi$ .  $\mathcal{A}$  will do the following.

- If  $Q_1 = \exists$ , then output 1 if and only if at least one of  $\mathcal{A}(\psi_{x_1=0})$  and  $\mathcal{A}(\psi_{x_1=1})$  output 1.
- If  $Q_1 = \forall$ , then output 1 if and only if both of  $\mathcal{A}(\psi_{x_1=0})$  and  $\mathcal{A}(\psi_{x_1=1})$  output 1.

It is clear that  $\mathcal{A}$  works correctly; we now show that  $\mathcal{A}$  can be implemented to use only polynomial space. Let  $s_{n,m}$  be the space used by  $\mathcal{A}$  on formulae with *n* variables and total size *m*. After computing  $\mathcal{A}(\psi_{x_1=0})$ , the algorithm can reuse the same space to compute  $\mathcal{A}(\psi_{x_1=1})$ , needing only to retain one bit to store the value of  $\mathcal{A}(\psi_{x_1=0})$ . At each step the algorithm also needs to use only  $O(m)$  space to store a description of  $\psi_{x_1=0}$  and  $\psi_{x_1=1}$ , to be passed to the next recursive call of  $\mathcal{A}$ . Thus  $s_{n,m} = s_{n-1,m} + O(m)$ , so  $s_{n,m} = O(nm)$ .

We now show that, for any  $\mathcal{L} \in \text{PSPACE}$ ,  $\mathcal{L} \leq_P \text{TQBF}$ . The proof will be based on the idea of configuration graphs which we saw previously. We are given a Turing machine *M* that decides  $\mathcal{L}$  in space  $s(n)$ . For each  $x \in \{0, 1\}^n$ , we aim to construct, in polynomial time, a quantified boolean formula  $\psi$  of size  $O(s(n)^2)$  which is true if and only if *M* accepts *x*. We will achieve this by producing a sequence of quantified formulae  $\psi_i$  such that  $\psi_i(a, b)$  is true if and only if there is a path of length at most  $2^i$  in  $G_{M,x}$  from *a* to *b*. We can then output  $\psi_{O(s(n))}(c_S, c_A)$ , where  $c_S, c_A$ ,

$c_A$  are the start and accepting configurations of  $M$ . As the configuration graph of  $M$  has at most  $2^{O(s(n))}$  nodes, this suffices to determine whether  $M$  accepts  $\mathcal{L}$ .

We first make the following claim: there is an  $O(s(n))$ -size CNF formula  $\phi$  such that, for any two configurations  $c, c'$ ,  $\phi(c, c') = 1$  if and only if there is an arc  $c \rightarrow c'$  in  $G_{M,x}$ . To see this, note that each configuration can be encoded as an  $O(s(n))$ -bit string and so, as in the proof of the Cook-Levin theorem, whether one configuration can follow from another can be encoded as the AND of at most  $O(s(n))$  constant-sized checks. This implies that we can construct  $\psi_0(a, b)$  by taking the OR of the formula  $\phi$  and the formula which checks whether  $a$  and  $b$  are equal.

We now want to use this idea to construct  $\psi_{i+1}$ , given  $\psi_i$ . The intuitive way of doing this would be to write down the quantified formula

$$\exists c \psi_i(a, c) \wedge \psi_i(c, b),$$

which encodes the claim that (in words) “there is a path from  $a$  to  $b$  of length at most  $2^{i+1}$  if and only if there exists some  $c$  such that there is a path from  $a$  to  $c$  of length at most  $2^i$ , and a path from  $c$  to  $b$  of length at most  $2^i$ ”. However, this approach would result in obtaining a formula of exponential size, as at each step  $\psi_i$  could double in size. We therefore instead define  $\psi_{i+1}(a, b)$  as

$$\exists c \forall x, y : ((x = a \wedge y = c) \vee (x = c \wedge y = b)) \Rightarrow \psi_i(x, y),$$

which encodes the same claim (as one can verify). Then we have that the size of  $\psi_{i+1}$  is at most the size of  $\psi_i$ , plus  $O(s(n))$ . Therefore,  $\psi_{s(n)}$  is of size  $O(s(n)^2)$ . It finally remains to convert the formula  $\psi_{O(s(n))}$  to normal form, which can be done in polynomial time.  $\square$

TQBF remains PSPACE-complete even if we restrict the boolean formula  $\phi$  to be in CNF. To see this, observe that for any  $x, y$ , “ $x \Rightarrow y$ ” can be written as  $\neg x \vee y$ , implying that for each  $i$ ,  $\psi_i$  can be written in DNF without significantly increasing in size. Taking the negation of the final formula  $\psi_{s(n)}$  gives a CNF formula for the *complement* of the PSPACE language  $\mathcal{L}$ ; but  $\text{co-PSPACE} = \text{PSPACE}$ , so the claim follows. The very observant reader may have noticed that the above proof that for any  $\mathcal{L} \in \text{PSPACE}$ ,  $\mathcal{L} \leq_P \text{TQBF}$  also works without change for nondeterministic Turing machines. This implies that (as  $\text{TQBF} \in \text{PSPACE}$ ) in fact  $\text{PSPACE} = \text{NPSPACE}$ ! This is in stark contrast to our intuition that  $\text{P} \neq \text{NP}$ .

## 8.2 Log space

The PSPACE vs. NPSPACE question has a log-space analogue: does  $\text{L} = \text{NL}$ ? Unlike PSPACE vs. NPSPACE, this question is open, although some interesting partial results towards resolving it are known. We begin by giving a natural complete problem for NL.

Given that  $\text{NL} \subseteq \text{P}$  by Theorem 8.1, we cannot use the same notion of polynomial-time reductions that we used before to sensibly define what it means for a problem to be NL-complete (which you will show as an exercise). The notion we will use instead is that of log-space reductions, the definition of which requires a little care.

We say that  $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$  is implicitly log-space computable if  $f$  is polynomially bounded (i.e. there exist  $b, c$  such that  $|f(x)| \leq b|x|^c$  for all  $x \in \{0, 1\}^*$ ), and the languages  $\mathcal{L}_f = \{(x, i) : f(x)_i = 1\}$  and  $\mathcal{L}'_f = \{(x, i) : |f(x)| \geq i\}$  are in L. That is, given  $x$ , we can determine any individual bit of  $f(x)$ , and the length of  $f(x)$ , using a log-space Turing machine.

We then say that  $\mathcal{A}$  is log-space reducible to  $\mathcal{B}$ , and write  $\mathcal{A} \leq_L \mathcal{B}$ , if there is an implicitly log-space computable function  $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$  such that  $x \in \mathcal{A}$  if and only if  $f(x) \in \mathcal{B}$ . We finally say that a language  $\mathcal{A}$  is NL-complete if  $\mathcal{A} \in \text{NL}$ , and for all  $\mathcal{B} \in \text{NL}$ ,  $\mathcal{B} \leq_L \mathcal{A}$ .

**Theorem 8.4.** *PATH is NL-complete.*

*Proof.* We first show that PATH is in NL. If there is a path from node  $s$  to node  $t$ , a nondeterministic machine can find it by repeatedly guessing the neighbour of the current node which is on the path. This only requires the number of steps taken on the path to be known, and the identity of the current node. As the length of the path can be taken to be at most  $n$ , we only need  $O(\log n)$  space to store this information.

Let  $\mathcal{L} \in \text{NL}$ ; we now show that  $\mathcal{L} \leq_L \text{PATH}$ . Let  $M$  be a log-space NDTM which decides  $\mathcal{L}$ , and for any input  $x$  of size  $n$ , consider the configuration graph  $G_{M,x}$  (assuming as before that  $M$  has starting configuration  $c_S$ , and a single accepting configuration  $c_A$ ). Our reduction will take the pair  $(M, x)$  as input and produce the adjacency matrix of  $G_{M,x}$ . As there is a path in this graph from  $c_S$  to  $c_A$  if and only if  $M$  accepts  $x$ , an algorithm for PATH can be used to determine from this adjacency matrix whether  $M$  accepts  $x$ . It remains to show that this reduction is implicitly log-space computable. This follows because, given a pair of configurations  $(c, c')$ , in space  $O(|c| + |c'|) = O(\log n)$  a Turing machine can examine each of them and determine whether  $c'$  is one of the two configurations which can follow  $c$ , based on  $M$ 's transition rules.  $\square$

By Theorem 8.4, resolving the  $L \stackrel{?}{=} \text{NL}$  question is equivalent to determining whether there is a log-space algorithm for PATH. While it is not known whether such an algorithm exists, it has recently been shown that PATH restricted to *undirected* graphs is indeed in L. We also have the following result.

**Theorem 8.5.**  $\text{PATH} \in \text{SPACE}(\log^2 n)$ .

*Proof.* The basic idea is very similar to the second part of the proof of Theorem 8.3. Let  $G$  be a directed graph with  $n$  nodes, and let  $\text{PATH}_i(a, b)$  be the following problem: given  $G$  (as an adjacency matrix) and nodes  $a$  and  $b$  in  $G$ , is there a path from  $a$  to  $b$  of length at most  $2^i$ ? Observe that if we can solve  $\text{PATH}_{\lceil \log_2 n \rceil}$ , we can solve PATH, as no path need be of length longer than  $n$ . We solve  $\text{PATH}_i$  by a recursive technique. If  $i = 0$ , we can solve  $\text{PATH}_0(a, b)$  simply by checking whether  $a = b$ , or there is an arc from  $a$  to  $b$ . If  $i \geq 1$ , we compute  $\text{PATH}_i(a, b)$  as follows: output true if and only if there exists a node  $c$  such that  $\text{PATH}_{i-1}(a, c)$  is true and  $\text{PATH}_{i-1}(c, b)$  is true. To see that this is correct, observe that (as in the proof of Theorem 8.3) there is a path of length  $2^i$  from  $a$  to  $b$  if and only if there is a path of length at most  $2^{i-1}$  from  $a$  to some ‘‘midpoint’’  $c$ , and a path of length at most  $2^{i-1}$  from  $c$  to  $b$ .

It remains to implement this computation space-efficiently. But this can be done using exactly the same idea as in the first part of the proof of Theorem 8.3; we obtain that the space used is  $O(\log^2 n)$ .  $\square$

Theorem 8.5 has the following corollary, a generalisation of the result that  $\text{PSPACE} = \text{NPSpace}$ .

**Theorem 8.6** (Savitch’s Theorem). *For any function  $s : \mathbb{N} \rightarrow \mathbb{N}$  such that  $s(n) \geq \log_2 n$ ,  $\text{NPSpace}(s(n)) \subseteq \text{SPACE}(s(n)^2)$ .*

*Proof.* For any  $\mathcal{L} \in \text{NPSpace}(s(n))$ , given an NDTM  $M$  which decides  $\mathcal{L}$  in space  $s(n)$ , we simply run the algorithm of Theorem 8.5 on the configuration graph of  $M$ , which as discussed before has at most  $2^{O(s(n))}$  nodes.  $\square$

Another question about space complexity which has been resolved and stands in contrast to the corresponding time complexity conjecture is the NL vs. co-NL question.

**Theorem 8.7** (Immerman-Szelepcsényi Theorem).  $\text{NL} = \text{co-NL}$ .

The technique used to prove Theorem 8.7 is a direct attack: giving a nondeterministic log-space algorithm for the complement of the PATH problem. It is worth pausing to note that it is somewhat surprising that this works; while it is easy to certify that there is a path between two nodes  $s$  and  $t$  (just give the path), it is less clear how to certify that there is *no* path from  $s$  to  $t$ .

*Proof of Theorem 8.7.* The complement of the PATH problem,  $\overline{\text{PATH}}$ , may be defined as follows. Given a graph  $G$  with  $n$  nodes, and specified nodes  $s, t$ , output 1 if there is no path from  $s$  to  $t$ , and 0 otherwise. We prove that  $\overline{\text{PATH}}$  is in NL by showing the existence of a certificate, which can be checked with read-once access in logarithmic space, that there is no path from  $s$  to  $t$ . The certificate will be built from a sequence of subsidiary certificates. Let  $C_i$  denote the set of nodes which can be reached from  $s$  in at most  $i$  steps, and write  $c_i = |C_i|$ ; we want to certify that  $t \notin C_n$ . We know (see the proof of Theorem 8.4) that, for each  $i$  and any  $v$ , there exists a certificate  $\text{Path}_i(s, v)$  which is read-once checkable in logarithmic space and certifies that there is a path of length at most  $i$  from  $s$  to  $v$ , i.e. proves that  $v \in C_i$ . We now use an inductive procedure to design two additional certificates.

1. A certificate that  $v \notin C_i$ , assuming that the verifier knows  $|C_i|$ .
2. A certificate that  $|C_i| = c$  for some  $c$ , assuming that the verifier knows  $|C_{i-1}|$ .

As the verifier already knows that  $C_0 = \{s\}$  at the beginning, using these two types of certificate for  $i = 1, \dots, n$  enables the verifier to determine that  $t \notin C_n$ . We now show how to construct them.

1. The certificate is simply the list of all nodes in  $C_i$ , together with their certificates that they are in  $C_i$ :

$$v_1, \text{Path}_i(s, v_1), \dots, v_{c_i}, \text{Path}_i(s, v_{c_i})$$

for  $v_1, \dots, v_{c_i} \in C_i$  in ascending order. The verifier checks that (1) the number of nodes is equal to  $c_i$ ; (2) the nodes are listed in ascending order; (3) none of the nodes is equal to  $v$ ; (4) each certificate is valid. It is easy to see that all of these checks can be done with read-once access to the whole certificate. If  $v \notin C_i$ , such a certificate exists; however, if  $v \in C_i$  no such list of nodes can convince the verifier.

In fact, a similar idea can be used to certify that  $v \notin C_i$ , given  $|C_{i-1}|$ . The only difference is that in step (3) the verifier checks that neither  $v$ , nor any neighbour of  $v$ , appears in the list of nodes in  $C_{i-1}$ . This procedure works because  $v \in C_i$  if and only if there exists  $u \in C_{i-1}$  such that either  $u = v$  or  $u \rightarrow v$ .

2. We have described certificates for each node  $v$  to prove that either  $v \in C_i$  or  $v \notin C_i$ . The certificate that  $|C_i| = c$  is just the list of these certificates for each node  $v$  in ascending order. The verifier simply checks all these certificates and accepts if and only if the number of nodes certified to be in  $C_i$  is equal to  $c$ .

□

### 8.3 Historical notes and further reading

Arora-Barak has a general discussion of space complexity in chapter 4. The Immerman-Szelepcsényi Theorem was proven independently by Immerman and Szelepcsényi in the late 80's. The fact that the PATH problem restricted to undirected graphs is in L was shown by Reingold in 2005; Arora-Barak chapter 21 includes a proof.

## 9 Randomised algorithms

The goal of the Turing machine model introduced in Section 3 is to encompass the notion of computation in the physical world. However, there is an aspect of the physical world which this model does not seem to capture: randomness. While it is a philosophical question whether true randomness “really” exists in the universe, for practical purposes it does appear that there exist devices (e.g. coins) which can be used to generate apparently random numbers, which could be used in computation. It is less clear whether such randomness is actually useful, so we begin by discussing two examples of non-trivial randomised algorithms.

First we will give some inequalities which will be very useful in the analysis of such algorithms, beginning by recalling two basic facts which will be used throughout: linearity of expectation ( $\mathbb{E}[X + Y] = \mathbb{E}[X] + \mathbb{E}[Y]$ ) and the union bound  $\Pr[A \cup B] \leq \Pr[A] + \Pr[B]$ .

**Theorem 9.1** (Markov’s inequality). *Let  $X$  be a random variable. Then, for any  $c > 0$ ,*

$$\Pr[|X| \geq c] \leq \frac{\mathbb{E}[|X|]}{c}.$$

*Proof.* Let  $p_x = \Pr[|X| = x]$ . Then, for any  $c > 0$ ,

$$\Pr[|X| \geq c] = \sum_{x \geq c} p_x \leq \frac{1}{c} \sum_{x \geq c} xp_x \leq \frac{1}{c} \sum_x xp_x = \frac{\mathbb{E}[|X|]}{c}.$$

□

A Bernoulli random variable is a random variable which takes the value 1 with probability  $p$ , for some  $0 \leq p \leq 1$ , otherwise taking the value 0. The sum of  $n$  independent Bernoulli random variables is tightly concentrated around its mean. This idea can be made precise using Chernoff bounds (also known as Chernoff-Hoeffding bounds or Hoeffding bounds), which are a fundamental tool in the analysis of randomised algorithms.

**Theorem 9.2** (Chernoff bound). *Let  $X = \sum_i X_i$ , where each  $X_i$  is an independent Bernoulli random variable taking the value 1 with probability  $p_i$ . Write  $\mu = \mathbb{E}[X] = \sum_i p_i$ . Then, for all  $\delta > 0$ ,*

$$\begin{aligned} \Pr[X \geq (1 + \delta)\mu] &\leq \left( \frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu, \text{ and} \\ \Pr[X \leq (1 - \delta)\mu] &\leq \left( \frac{e^{-\delta}}{(1 - \delta)^{1-\delta}} \right)^\mu. \end{aligned}$$

*Proof.* The basic idea is to apply Markov’s inequality to the random variable  $e^{tX}$ , for some  $t$ , and then to optimise over  $t$ . For the first part, we have the following chain of equalities and inequalities.

$$\begin{aligned}
\Pr \left[ \sum_i X_i \geq (1 + \delta)\mu \right] &= \Pr \left[ e^{t \sum_i X_i} \geq e^{t(1+\delta)\mu} \right] && \text{[for any } t > 0] \\
&\leq \frac{\mathbb{E} \left[ e^{t \sum_i X_i} \right]}{e^{t(1+\delta)\mu}} && \text{[Markov's inequality]} \\
&= \frac{\prod_i \mathbb{E} \left[ e^{t X_i} \right]}{e^{t(1+\delta)\mu}} && \text{[by independence of } X_i \text{'s]} \\
&= \frac{\prod_i (1 + p_i(e^t - 1))}{e^{t(1+\delta)\mu}} && \text{[by definition of } X_i \text{'s]} \\
&\leq \frac{e^{(e^t - 1) \sum_i p_i}}{e^{t(1+\delta)\mu}} && \text{[using the inequality } 1 + x \leq e^x] \\
&= e^{\mu((e^t - 1) - t(1+\delta))}.
\end{aligned}$$

We now minimise this expression over  $t$ . The minimum is obtained at  $t = \ln(1 + \delta)$ , which gives

$$\Pr \left[ \sum_i X_i \geq (1 + \delta)\mu \right] \leq \left( \frac{e^\delta}{(1 + \delta)^{1+\delta}} \right)^\mu$$

as claimed. For the second part, we follow exactly the same procedure but start with

$$\Pr \left[ \sum_i X_i \leq (1 - \delta)\mu \right] = \Pr \left[ - \sum_i X_i \geq -(1 - \delta)\mu \right].$$

□

One can simplify these bounds to obtain the following corollary.

**Corollary 9.3.** *Let  $X = \sum_i X_i$ , where each  $X_i$  is an independent Bernoulli random variable taking the value 1 with probability  $p_i$ . Write  $\mu = \mathbb{E}[X] = \sum_i p_i$ . Then, for all  $0 \leq \delta \leq 1$ ,*

$$\begin{aligned}
\Pr[X \geq (1 + \delta)\mu] &\leq e^{-\mu\delta^2/3}, \text{ and} \\
\Pr[X \leq (1 - \delta)\mu] &\leq e^{-\mu\delta^2/2}.
\end{aligned}$$

*Proof.* For the first part, we examine the Taylor expansion of  $\ln(1 + \delta)$  to obtain

$$(1 + \delta) \ln(1 + \delta) = \delta + \frac{\delta^2}{2} - \frac{\delta^3}{6} + \frac{\delta^4}{12} - \dots \geq \delta + \frac{\delta^2}{2} - \frac{\delta^3}{6} \geq \delta + \frac{\delta^2}{3},$$

valid for  $0 \leq \delta \leq 1$ , which implies  $(1 + \delta)^{1+\delta} \geq e^{\delta + \delta^2/3}$ . For the second part, we bound  $(1 - \delta)^{1-\delta} \geq e^{-\delta + \delta^2/2}$ , which similarly follows from the Taylor expansion of  $\ln(1 - \delta)$ . □

Finally, an additive form of these bounds will often be useful, which we state as the following immediate corollary (with suboptimal constants).

**Corollary 9.4.** *Let  $\bar{X} = \frac{1}{n} \sum_{i=1}^n X_i$ , where each  $X_i$  is an independent Bernoulli random variable taking the value 1 with probability  $p_i$ . Write  $\mu = \mathbb{E}[\bar{X}] = \frac{1}{n} \sum_{i=1}^n p_i$ . Then, for all  $0 \leq \delta \leq 1$ ,*

$$\begin{aligned}
\Pr[\bar{X} \geq \mu + \delta] &\leq e^{-n\delta^2/3}, \text{ and} \\
\Pr[\bar{X} \leq \mu - \delta] &\leq e^{-n\delta^2/2}.
\end{aligned}$$

*Proof.* Immediate from Corollary 9.3, observing that  $0 \leq \mu \leq 1$ . □

Observe that the phrasing of Corollary 9.4 highlights the fact that, as  $n$  increases,  $\bar{X}$  becomes rapidly more concentrated around  $\mu$ .

## 9.1 Primality testing

Imagine that we are given an  $n$ -digit integer  $N$ , and would like to determine whether  $N$  is prime. The input size is  $n$ , so we would like an algorithm that runs in time  $\text{poly}(n)$ , i.e.  $\text{poly}(\log N)$ . In particular, note that the naïve algorithm which tests all potential divisors of  $N$  (from 2 up to  $\lfloor \sqrt{N} \rfloor$ ) does not run in time  $\text{poly}(n)$ . We now give a polynomial-time randomised test for primality, which is known as the Miller-Rabin test.

1. Assume that  $N$  is odd, and write  $N - 1$  as  $2^j s$  for some odd  $s$  and integer  $j$ .
2. Pick a random integer  $r \in \{1, \dots, N - 1\}$ .
3. Compute the sequence  $S = r^s, r^{2s}, r^{4s}, \dots, r^{2^j s}$  modulo  $N$  by repeated squaring.
4. Output “composite” if either:
  - (a)  $r^{2^j s} \not\equiv 1$ , or
  - (b) for some  $i$ , the sequence goes from  $r^{2^i s} \notin \{1, N - 1\}$  to  $r^{2^{i+1} s} \equiv 1$ .
5. Otherwise, output “prime”.

**Theorem 9.5.** *If  $N$  is prime, the Miller-Rabin test always returns “prime”. If  $N$  is composite, the Miller-Rabin test returns “composite” with probability at least  $1/2$ .*

For any  $n$ , let  $\mathbb{Z}_n$  denote the additive group of integers modulo  $n$ , and let  $\mathbb{Z}_n^\times$  denote the multiplicative group of integers modulo  $n$ , i.e. the group consisting of integers  $a$  such that  $0 < a < n$  and  $a$  is coprime to  $n$ , under multiplication. Recall that  $a$  is coprime to  $n$  if and only if there exists an integer  $b$  such that  $ab \equiv 1 \pmod{n}$ . In order to prove Theorem 9.5, we will need the following two facts from basic number theory, which we state without proof.

- Fermat’s little theorem: For any prime  $p$ , and any positive integer  $a < p$ ,  $a^{p-1} \equiv 1 \pmod{p}$ .
- The Chinese remainder theorem: For a positive integer  $N$  with prime factorisation  $N = p_1^{r_1} \dots p_k^{r_k}$ ,  $\mathbb{Z}_N^\times \cong \mathbb{Z}_{p_1^{r_1}}^\times \times \dots \times \mathbb{Z}_{p_k^{r_k}}^\times$ .

*Proof of Theorem 9.5.* First, if  $N$  is prime, then by Fermat’s little theorem  $r^{2^j s} \equiv 1 \pmod{N}$ , so the last element of  $S$  is 1. Therefore, the sequence is either all 1’s, or becomes 1 following a squaring operation. But in the field  $\mathbb{Z}_N$  the only roots of the equation  $x^2 = 1$  are  $x = \pm 1$ , so the test always outputs “prime” when  $N$  is prime. We now use some group theory to prove that the test works with probability at least  $1/2$  when  $N$  is composite.

First assume that there exists an  $a \in \mathbb{Z}_N^\times$  such that  $a^{2^j s} \not\equiv 1$  (i.e. via Fermat’s little theorem,  $a$  is a “witness” that  $N$  is composite). The  $a$ ’s such that  $a^{2^j s} \equiv 1 \pmod{N}$  form a subgroup of  $\mathbb{Z}_N^\times$ . As this subgroup is proper, it contains at most half of the elements of  $\mathbb{Z}_N^\times$ . So, when  $r$  is picked at random, it is either coprime to  $N$  (in which case with probability at least  $1/2$ ,  $r^{2^j s} \not\equiv 1$ ), or it is not coprime to  $N$  (in which case  $r^{2^j s} \not\equiv 1$  always).

What if this does not work – i.e. for all  $a \in \mathbb{Z}_N^\times$ ,  $a^{2^j s} \equiv 1$ ? Such  $N$  exist and are called *Carmichael numbers* (the smallest is 561). In this case, we show that the algorithm still works. Consider the map  $r \mapsto r^s$  which we perform first. This is a homomorphism from the group  $\mathbb{Z}_N^\times$  to some subgroup  $H_1$  (i.e.  $(ab)^s = a^s b^s$  for all  $a, b \in \mathbb{Z}_N^\times$ ), so  $r^s$  is uniformly distributed in  $H_1$ . We then repeatedly apply the map  $a \mapsto a^2$ , giving a sequence of subgroups  $H_2, \dots, H_{j+1}$ , where  $H_{j+1} = \{1\}$  because  $N$  is a Carmichael number. This implies that the order of every element of each subgroup  $H_i$  is a power of 2, so  $|H_i|$  is a power of 2.

Now consider the subgroup  $I$  which is the last subgroup before we get to  $\{1\}$ . The numbers  $r$  that pass the Miller-Rabin test must have been mapped to either 1 or  $-1$  at this stage.  $I$  may or may not contain the element  $-1$ . If it does not contain  $-1$ , as  $|I| \geq 2$ ,  $r$  fails the Miller-Rabin test with probability at least  $1/2$ . If it contains  $-1$  and at least one more element (other than 1),  $|I| \geq 4$  (as  $|I|$  is a power of 2), so the same conclusion follows.

We now show that if  $-1 \in I$ , so is another element (other than 1). If  $N$  is not a power of a prime, by the Chinese remainder theorem  $\mathbb{Z}_N^\times$  is isomorphic to a product  $J \times K$  of nontrivial groups. So the element  $-1 \in \mathbb{Z}_N^\times$  can be represented as a pair  $(-1_J, -1_K)$ . Consider the elements  $(-1_J, 1_K)$  and  $(1_J, -1_K)$ . If  $-1 \in I$ , this is because the map  $r \mapsto r^{2^i s}$ , for some  $i$ , can be applied to some element  $(x, y) \in \mathbb{Z}_N^\times$  to obtain  $(-1_J, -1_K)$ . But then the same map can be applied to  $(x, 1)$  and  $(1, y)$  to obtain  $(-1_J, 1_K)$ ,  $(1_J, -1_K)$ , so these two elements are also in  $I$ .

We finally claim that, if  $N$  is a power of a prime, it is not a Carmichael number, so the test also works in this case. Assume  $N = p^k$  and take  $a = p + 1$ . We show that  $a^N \not\equiv a \pmod{p^2}$ , which implies that  $a^N \not\equiv a \pmod{p^k}$ , so  $a^{N-1} \not\equiv 1 \pmod{N}$  and hence  $N$  is not a Carmichael number. By taking  $(p + 1)^p = \sum_{\ell=0}^p \binom{p}{\ell} p^\ell \equiv 1 \pmod{p^2}$ , we have  $a^N \equiv 1 \pmod{p^2}$ , implying  $a^N \not\equiv a$ .  $\square$

It has been shown quite recently that in fact primality testing is in P. However, the best deterministic algorithm currently known still has a fairly large runtime ( $O(n^6 \text{polylog}(n))$ ), so in practice randomised algorithms are still used for primality testing.

## 9.2 A randomised algorithm for $k$ -SAT

We now return once more to the much-studied  $k$ -SAT problem (recall: given a boolean formula  $\phi(x_1, \dots, x_n)$  in CNF with at most  $k$  variables per clause, does there exist an assignment to the variables  $x_i$  such that  $\phi$  evaluates to true?). Consider the following algorithm  $\mathcal{A}$ .

1. Initially assign  $x_i$  randomly for all  $i$ .
2. Repeat the following  $t$  times, for some  $t$  to be determined:
  - (a) If  $\phi$  is satisfied by  $x$ , return that  $\phi$  is satisfiable.
  - (b) Otherwise, pick an arbitrary clause  $C$  in  $\phi$  which is not satisfied by  $x$ . Pick one of the variables  $x_i$  in  $C$  at random and update  $x$  by mapping  $x_i \mapsto \neg x_i$ .
3. Return that  $\phi$  is not satisfiable.

We will prove the following claims:

**Theorem 9.6.** *If  $k = 2$  and we take  $t = 2n^2$ ,  $\mathcal{A}$  succeeds with probability at least  $1/2$ .*

**Theorem 9.7.** *If  $k > 2$  and we take  $t = 3n$ ,  $\mathcal{A}$  succeeds with probability at least  $\left(\frac{k}{2(k-1)}\right)^n / \text{poly}(n)$ .*






Figure 12: The random walk performed by algorithm  $\mathcal{A}$  for  $k$ -SAT.

Observe that, for any integer  $K > 0$ , Theorem 9.6 implies an algorithm for 2-SAT which fails with probability at most  $2^{-K}$  and uses time  $O(Kn^2)$ : simply repeat  $\mathcal{A}$   $K$  times, return “satisfiable” if any of these uses of  $\mathcal{A}$  returned a satisfying assignment, and return “unsatisfiable” otherwise. The probability that  $\mathcal{A}$  fails every time is at most  $2^{-K}$ . Similarly, Theorem 9.7 implies an algorithm for  $k$ -SAT which succeeds with probability at least 0.99 (say) in time  $O\left(\left(\frac{2(k-1)}{k}\right)^n \text{poly}(n)\right)$ . For small  $k$ , this can be much better than the trivial algorithm of just trying all possible assignments to  $x$ ; e.g. for  $k = 3$  we get  $O((4/3)^n \text{poly}(n))$ , as compared with the  $\Omega(2^n)$  time taken by the trivial algorithm. This illustrates that even NP-complete problems can have non-trivial algorithms.

The proofs of Theorems 9.6 and 9.7 will both be based on the same idea, of reducing the algorithm’s behaviour to a random walk. If  $\phi$  is satisfiable, there exists a satisfying assignment  $x^*$ . For any  $x \in \{0, 1\}^n$ , let  $d(x, x^*)$  be the Hamming distance between  $x$  and  $x^*$ , i.e. the number of positions at which  $x$  and  $x^*$  differ. At each step of the algorithm, as a single bit of  $x$  is flipped,  $d(x, x^*)$  either increases by 1 or decreases by 1. If  $d(x, x^*)$  ever becomes 0, the algorithm has found a satisfying assignment to  $\phi$  (i.e.  $x^*$ ). Importantly, the probability that  $d(x, x^*)$  decreases by 1 is at least  $1/k$ , which we see as follows. Let  $C$  be the clause picked by the algorithm, and assume it contains variables  $x_{i_1}, \dots, x_{i_k}$ . As  $x^*$  satisfies  $\phi$  (and hence  $C$ ) but  $x$  does not satisfy  $C$ ,  $x$  and  $x^*$  must differ in at least one position  $i \in \{i_1, \dots, i_k\}$ . Hence the probability that  $i$  is picked by the algorithm is at least  $1/k$ . Therefore, the behaviour of the algorithm can be mapped to a random walk on the line indexed by integers  $0, \dots, n$ , which moves left with probability at least  $1/k$  at each step; see Figure 12 for an illustration. We simplify this to a random walk which moves left with probability exactly  $1/k$  at each step; calculating the expected number of steps until this walk hits 0 then gives an upper bound on the expected number of steps until  $\mathcal{A}$  finds  $x^*$ . We have assumed that each clause  $C$  contains exactly  $k$  variables, but if  $C$  contains fewer than  $k$  variables this can only increase the probability of moving left.

We are now ready to prove Theorems 9.6 and 9.7.

*Proof of Theorem 9.6.* The claim will follow from Markov’s inequality by showing that the expected time to hit 0 is at most  $n^2$ . This is a well-known result from probability theory, but we give a simple direct argument.

For  $0 \leq i \leq n$ , let  $T(i)$  denote the expected number of steps until the random walk hits 0, given that it started at position  $i$ . It should be clear that:

- $T(0) = 0$ ;
- $T(n) = 1 + T(n - 1)$  [the walk cannot move further right than  $n$ ];
- $T(i) = 1 + \frac{1}{2}(T(i + 1) + T(i - 1))$  for  $1 \leq i \leq n - 1$ .

This recurrence can readily be solved to yield  $T(i) = i(2n - i)$ . Thus the expected time to hit 0 is at most  $n^2$  for any initial assignment  $x$ .  $\square$

*Proof of Theorem 9.7.* The proof of this theorem is along the same lines as Theorem 9.6, but somewhat more complicated.

Imagine the walker starts at position  $j$  and walks for  $T$  steps. Let  $P_j$  be the probability to reach 0 at some time between 1 and  $T$ . For any integer  $\ell$ , the probability to go left  $j + \ell$  times during the walk is lower bounded by

$$\binom{T}{j + \ell} k^{-(j+\ell)} (1 - 1/k)^{T-(j+\ell)}.$$

Now take  $T = (1 + 2\alpha)j$ ,  $\ell = \alpha j$ , for some  $\alpha$ . The probability to go left  $(1 + \alpha)j$  times (and hence right  $\alpha j$  times) lower bounds the probability to end up at 0 after at most  $T$  steps, giving the bound that

$$P_j \geq \binom{(1 + 2\alpha)j}{\alpha j} k^{-(1+\alpha)j} (1 - 1/k)^{(1-\alpha)j}.$$

We now choose  $\alpha = 1/(k - 2)$ . Observe that this implies that  $T \leq (1 + 2/(k - 2))n \leq 3n$ . Up to  $\text{poly}(j)$  terms, we have

$$\binom{(1 + 2\alpha)j}{\alpha j} = \left( \left( \frac{1 + 2\alpha}{\alpha} \right)^\alpha \left( \frac{1 + 2\alpha}{1 + \alpha} \right)^{1+\alpha} \right)^j,$$

which follows from Stirling's approximation

$$\sqrt{2\pi n} \left( \frac{n}{e} \right)^n e^{1/(12n+1)} \leq n! \leq \sqrt{2\pi n} \left( \frac{n}{e} \right)^n e^{1/(12n)},$$

so up to  $\text{poly}(j)$  terms,

$$P_j \geq \left( \left( \frac{1 + 2\alpha}{\alpha} \right)^\alpha \left( \frac{1 + 2\alpha}{1 + \alpha} \right)^{1+\alpha} \right)^j k^{-(1+\alpha)j} (1 - 1/k)^{(1-\alpha)j} = \left( \frac{1}{k-1} \right)^j.$$

The start position  $j$  is given by the Hamming weight (i.e. the number of 1s) of a random  $n$ -bit string  $x$ . Thus, using the binomial theorem, up to polynomial terms the probability to reach 0 in  $T \leq 3n$  steps is equal to

$$P = \frac{1}{2^n} \sum_{j=0}^n \binom{n}{j} P_j \geq \frac{1}{2^n} \sum_{j=0}^n \binom{n}{j} \left( \frac{1}{k-1} \right)^j = \left( \frac{k}{2(k-1)} \right)^n$$

as claimed. □

### 9.3 Adding randomness to Turing machines

We have seen that randomness can apparently be a useful resource. We now formalise the ability to use randomness in algorithms via the model of the probabilistic Turing machine (PTM).

A PTM is a Turing machine  $M$  with two transition functions  $\delta_0, \delta_1$ . At each step of its execution, with probability  $1/2$   $M$  applies  $\delta_0$ , and with probability  $1/2$  it applies  $\delta_1$ . The machine only has two possible outputs: 1 (accept) or 0 (reject). For some function  $T : \mathbb{N} \rightarrow \mathbb{N}$ , we say that  $M$  halts in  $T(n)$  time if, for any input  $x$ ,  $M$  halts within  $T(|x|)$  steps (whatever random choices it makes). Let  $M(x)$  be the random variable denoting the output of  $M$  on input  $x$ .

We can now define the complexity class **BPP**, which encapsulates our idea of efficient probabilistic computation. For  $\mathcal{L} \subseteq \{0, 1\}^*$ , we say that  $M$  decides  $\mathcal{L}$  in time  $T(n)$  if, for every  $x \in \{0, 1\}^*$ ,  $M$  halts in  $T(|x|)$  steps (regardless of its random choices), and:

- For all  $x \in \mathcal{L}$ ,  $\Pr[M(x) = 1] \geq 2/3$ ;
- For all  $x \notin \mathcal{L}$ ,  $\Pr[M(x) = 1] \leq 1/3$ .

Let  $\text{BPTIME}(T(n))$  be the class of languages decided by some PTM in time  $O(T(n))$  and write

$$\text{BPP} = \bigcup_{c>0} \text{BPTIME}(n^c).$$

Observe that an equivalent alternative definition of PTMs would be to take a standard deterministic Turing machine  $M$ , and give  $M$  an additional “randomness” tape containing random bits.

The constants  $2/3$ ,  $1/3$  appearing in this definition may appear somewhat arbitrary. We now show that, in fact, they can be replaced with any pair of numbers either side of  $1/2$  separated by at least an inverse polynomial gap.

**Lemma 9.8.** *Assume there exists a PTM  $M$  that decides  $\mathcal{L} \subseteq \{0,1\}^*$  with success probability  $1/2 + \gamma$  using time  $T(n)$ , i.e.:*

- for all  $x \in \mathcal{L}$ ,  $\Pr[M(x) = 1] \geq 1/2 + \gamma$ , and
- for all  $x \notin \mathcal{L}$ ,  $\Pr[M(x) = 1] \leq 1/2 - \gamma$ .

Then  $\mathcal{L} \in \text{BPTIME}(T(n)/\gamma^2)$ .

*Proof.* We will boost the success probability of  $M$  using majority voting. Our algorithm is simply to run  $M$   $K$  times (for some  $K$  to be determined), obtaining outcomes  $y_1, \dots, y_K$ . If  $|\{i : y_i = 1\}| \geq K/2$ , we output 1, otherwise we output 0.

Assume  $x \notin \mathcal{L}$  (the case  $x \in \mathcal{L}$  is similar). Then each  $y_i \in \{0,1\}$  is an independent random variable with  $\mathbb{E}[y_i] \leq 1/2 - \gamma$ . Our algorithm will fail if  $\bar{Y} := \frac{1}{K} \sum_{i=1}^K y_i \geq 1/2$ . We have  $\mathbb{E}[\bar{Y}] \leq 1/2 - \gamma$ . By the Chernoff bound (Corollary 9.4),

$$\Pr[\bar{Y} \geq 1/2] \leq e^{-\gamma^2 K/3},$$

so taking  $K = O(1/\gamma^2)$  suffices to upper bound the probability of failure by  $1/3$ . Observe that, if we increase  $K$  further than this, the success probability goes to 1 exponentially fast!  $\square$

As we see from this proof, given a BPP algorithm for deciding some language  $\mathcal{L}$ , we can write down another BPP algorithm which decides  $\mathcal{L}$  with *exponentially small* worst-case failure probability. One may wonder whether it is possible to relax the constraints on the probabilities further without changing the definition of BPP. Let PP be the class of languages  $\mathcal{L}$  decided by a polynomial-time PTM such that

- For all  $x \in \mathcal{L}$ ,  $\Pr[M(x) = 1] > 1/2$ ;
- For all  $x \notin \mathcal{L}$ ,  $\Pr[M(x) = 1] \leq 1/2$ ;

Then we have the following result.

**Theorem 9.9.**  $\text{NP} \subseteq \text{PP}$ .

*Proof.* Suppose that  $\mathcal{L} \in \text{NP}$  and let  $M$  be a polynomial-time NDTM with two transition functions (i.e. at each computational step, two possibilities for what to do next). We can therefore treat  $M$  as a PTM by flipping a coin at each step to decide which transition function to use. Consider a machine  $M'$  which has an additional step at the beginning: with probability  $1/2$ , run  $M$ ; and with probability  $1/2$ , accept. If the input  $x \in \mathcal{L}$ , then  $M$  has at least one accepting path on  $x$ , so  $\Pr[M'(x) = 1] > 1/2$ . If  $x \notin \mathcal{L}$ , then  $M$  has no accepting paths, so  $\Pr[M'(x) = 1] = 1/2$ . Thus  $\mathcal{L} \in \text{PP}$ .  $\square$

It therefore seems unlikely that  $\text{BPP} = \text{PP}$ . Indeed, many conjecture that in fact  $\text{BPP} = \text{P}$ . This would accord with our intuition that, in order to simulate “truly random” numbers, it suffices to produce suitably good “pseudorandom” numbers deterministically. However, it is currently not even known whether  $\text{BPP} \subseteq \text{NP}$  (although it *is* known that  $\text{BPP} \subseteq \Sigma_2$ !). You will prove the weaker statement that  $\text{BPP} \subseteq \text{PSPACE}$  (indeed, even  $\text{PP} \subseteq \text{PSPACE}$ ) as an exercise.

## 9.4 Historical notes and further reading

There are many good sets of lecture notes online about randomised algorithms (the discussion here of primality testing is partly based on one such lecture<sup>1</sup>). Papadimitriou chapter 11 gives an extensive analysis of an alternative test for primality by Solovay and Strassen. The deterministic primality test was invented by Agrawal, Kayal and Saxena (two of whom were undergraduates at the time) in 2002.

The randomised algorithm for  $k$ -SAT given here was invented by Schöning in 1999; the case of 2-SAT was already known (Papadimitriou chapter 11 has a discussion). There are several other non-trivial (but still exponential-time!) algorithms known for  $k$ -SAT, some of which are deterministic.

The question of whether  $\text{BPP} = \text{P}$ , and “derandomisation” questions more generally, are discussed at length in Arora-Barak chapters 20 and 21.

---

<sup>1</sup><http://www.cse.buffalo.edu/~regan/cse681/notes/lectureB12.pdf>

## 10 Counting complexity and the class #P

Some of the complexity classes we have seen so far have nice definitions in terms of computational paths of NDTMs. In particular, we can write down the following characterisations.

- A language  $\mathcal{L} \in \mathsf{P}$  if there exists a polynomial-time NDTM  $M$  such that, on input  $x$ :
  1. For all  $x \in \mathcal{L}$ , all of  $M$ 's computational paths accept;
  2. For all  $x \notin \mathcal{L}$ , none of  $M$ 's computational paths accept.
- A language  $\mathcal{L} \in \mathsf{NP}$  if there exists a polynomial-time NDTM  $M$  such that, on input  $x$ :
  1. For all  $x \in \mathcal{L}$ , at least one of  $M$ 's computational paths accepts;
  2. For all  $x \notin \mathcal{L}$ , none of  $M$ 's computational paths accept.
- A language  $\mathcal{L} \in \mathsf{co-NP}$  if there exists a polynomial-time NDTM  $M$  such that, on input  $x$ :
  1. For all  $x \in \mathcal{L}$ , all of  $M$ 's computational paths accept;
  2. For all  $x \notin \mathcal{L}$ , at least one of  $M$ 's computational paths does not accept.
- A language  $\mathcal{L} \in \mathsf{BPP}$  if there exists a polynomial-time NDTM  $M$  such that, on input  $x$ :
  1. For all  $x \in \mathcal{L}$ , at least  $2/3$  of  $M$ 's computational paths accept;
  2. For all  $x \notin \mathcal{L}$ , at most  $1/3$  of  $M$ 's computational paths accept.
- A language  $\mathcal{L} \in \mathsf{PP}$  if there exists a polynomial-time NDTM  $M$  such that, on input  $x$ :
  1. For all  $x \in \mathcal{L}$ ,  $> 1/2$  of  $M$ 's computational paths accept;
  2. For all  $x \notin \mathcal{L}$ ,  $\leq 1/2$  of  $M$ 's computational paths accept.

A problem at least as difficult as all of the above complexity classes is calculating the number of accepting paths of an NDTM  $M$  on input  $x$ . This motivates the introduction of the complexity class #P (pronounced “sharp P”), which is defined as the class of all functions  $f$  such that  $f(x)$  is equal to the number of accepting paths of some NDTM  $M$  on input  $x$ . Note that this is a class of functional problems, rather than decision problems. Once again, we can equivalently define this class in terms of certificates, as follows:

**Definition 10.1.** A function  $f : \{0, 1\}^* \rightarrow \mathbb{N}$  is in #P if there exists a polynomial  $p : \mathbb{N} \rightarrow \mathbb{N}$  and a polynomial-time Turing machine  $M$  such that, for all  $x \in \{0, 1\}^*$ ,

$$f(x) = \left| \{y \in \{0, 1\}^{p(|x|)} : M(x, y) = 1\} \right|.$$

Informally speaking, #P is the class of problems which involve determining the size of a set (which is possibly exponentially large), such that membership in the set can be checked efficiently. It is perhaps easiest to understand this class by considering some simple examples of problems in #P.

- #SAT: given a boolean formula  $\phi$ , calculate the number of satisfying assignments of  $\phi$ . This is clearly at least as hard as SAT and TAUTOLOGY.

- #PATH: given a graph  $G$  and two vertices  $s$  and  $t$ , count the number of simple paths from  $s$  to  $t$  (a path is called simple if it does not visit any vertex twice).

As with NP, #P has a notion of complete problems, which are the “hardest” problems in #P. Recalling that FP is the class of polynomial-time computable functions, we define  $\text{FP}^f$  (for some function  $f : \{0,1\}^* \rightarrow \{0,1\}^*$ ) to be the class of functions which are computable in polynomial time, given access to an oracle for  $f$ . Generalising the definition of oracles for decision problems, we say that a Turing machine  $M$  has an oracle for  $f$  if it has an oracle for the language  $\mathcal{L}_f = \{(x, i) : f(x)_i = 1\}$ . In other words,  $M$  can extract arbitrary bits of  $f(x)$  “for free”.

**Definition 10.2.**  $f$  is #P-complete if  $f \in \text{\#P}$  and every function  $g \in \text{\#P}$  is in  $\text{FP}^f$ . Thus, if  $f$  is #P-complete and  $f \in \text{FP}$ ,  $\text{FP} = \text{\#P}$ .

Perhaps unsurprisingly, we have the following result.

**Theorem 10.1.** #SAT is #P-complete.

*Proof.* Imagine that, given the description of some NDTM  $M$ , we could output a boolean formula  $\phi$  such that the number of satisfying assignments of  $\phi$  is equal to the number of accepting paths of  $M$ . If this were the case, the theorem would be proven, as given an oracle for #SAT we could compute the number of accepting paths of  $M$ , for any NDTM  $M$ . Inspecting the proof of the Cook-Levin Theorem (Section 5.5) shows that this almost gives us such a reduction. All that is required is to constrain  $M$  to accept only after a fixed number of steps, and in a fixed configuration; both of these modifications are straightforward.  $\square$

For languages  $\mathcal{A}, \mathcal{B} \in \text{NP}$ , we say that a reduction from language  $\mathcal{A}$  to  $\mathcal{B}$  is *parsimonious* if it preserves the number of accepting paths, i.e. the number of solutions. The above theorem gives a parsimonious reduction from any language in NP to SAT. In fact, there also exists a parsimonious reduction from any language in NP to 3-SAT. We have technically not proven this yet, as the reduction from SAT to 3-SAT we previously gave is not parsimonious; we defer the proof of this claim until later, when we will give an alternative proof of the Cook-Levin theorem.

In order to prove that some counting problem in #P with underlying language  $\mathcal{L} \in \text{NP}$  is #P-complete, it therefore suffices to give a parsimonious reduction from 3-SAT to  $\mathcal{L}$ . More generally, it suffices to give a reduction which is parsimonious up to some easily computable additive or multiplicative constant; we also call reductions of this form parsimonious.

Although it is arguably unexciting that the counting variant of an NP-complete problem such as 3-SAT should be hard, it is more surprising that there exist #P-complete counting problems whose decision variants are in P. A simple such example is the problem of counting satisfying assignments to a *monotone* boolean formula, which is a boolean formula with no negations. For example,

$$\phi = (x_1 \wedge (x_2 \vee x_3 \vee x_4)) \vee (x_1 \vee x_3)$$

is a monotone boolean formula. Such formulae always have a satisfying assignment where all variables are set to true, so determining whether a monotone boolean formula has a satisfying assignment is clearly in P. However, we have the following result.

**Theorem 10.2.** The problem #MON of counting the number of satisfying assignments to a monotone formula is #P-complete.

*Proof.* #MON is clearly in #P. To prove #P-hardness, we give a reduction from #SAT to #MON. Given a boolean formula  $\phi(x_1, \dots, x_n)$ , we produce a new formula  $\psi(x_1, \dots, x_n, y_1, \dots, y_n)$  where  $\psi$  is the same as  $\phi$ , but with each occurrence of  $\neg x_i$  replaced with  $y_i$ , for all  $i$ . Write  $\#\phi$  for the number of satisfying assignments of  $\phi$ , for any boolean formula  $\phi$ . Then

$$\#\phi = \# \left( \psi \wedge \left( \bigwedge_{i=1}^n (x_i \vee y_i) \right) \wedge \left( \bigwedge_{i=1}^n \neg(x_i \wedge y_i) \right) \right);$$

the number of satisfying assignments of  $\phi$  is the same as the number of satisfying assignments of  $\psi$  where  $x_i \neq y_i$  for all  $i$ . Now write

$$\psi_A = \psi \wedge \left( \bigwedge_{i=1}^n (x_i \vee y_i) \right), \quad \psi_B = \left( \bigvee_{i=1}^n (x_i \wedge y_i) \right).$$

Then (using De Morgan's law)  $\#\phi = \#(\psi_A \wedge \neg\psi_B)$ , where  $\psi_A$  and  $\psi_B$  are both monotone formulae. But

$$\#(\psi_A \wedge \neg\psi_B) = \#\psi_A - \#(\psi_A \wedge \psi_B),$$

simply by counting the number of assignments that satisfy  $\psi_A$  but not  $\psi_B$ . Hence  $\#\phi$  can be computed as the difference between the numbers of satisfying assignments of two monotone formulae.  $\square$

Many important combinatorial problems also turn out to be #P-complete. An example, which we state without proof, is a problem from statistical physics. The *Ising model* was originally introduced as a model for studying ferromagnetism. An instance of the model is a set of  $n$  sites, each pair of which has an interaction energy  $V_{ij}$ . A configuration is an assignment  $\sigma \in \{\pm 1\}^n$ , where  $\sigma_i$  is the spin assigned to the  $i$ 'th site. The energy of a configuration is given by

$$H(\sigma) = - \sum_{i < j=1}^n V_{ij} \sigma_i \sigma_j - B \sum_{k=1}^n \sigma_k,$$

where  $B$  is some fixed magnetic field intensity. The central problem associated with the Ising model is to compute the partition function

$$Z = \sum_{\sigma \in \{\pm 1\}^n} e^{-\beta H(\sigma)}$$

for some  $\beta$ . Observe that  $Z$  is a sum over exponentially many terms, each of which can easily be computed. It turns out that computing  $Z$  is in fact #P-hard.

## 10.1 The permanent is #P-complete

We now give a natural mathematical problem which is also #P-complete. Let  $M$  be an  $n \times n$  matrix of integers. We say that  $M$  is a 0-1 matrix if it only contains 0's and 1's. The determinant of  $M$  can be defined as

$$\det(M) = \sum_{\sigma \in S_n} \text{sgn}(\sigma) \prod_{i=1}^n M_{i, \sigma(i)},$$

where the sum is over all permutations  $\sigma$  of  $\{1, \dots, n\}$ , and  $\text{sgn}(\sigma)$  is the sign of  $\sigma$ , i.e.  $\text{sgn}(\sigma) = 1$  if  $\sigma$  can be written as a product of an even number of transpositions, and  $\text{sgn}(\sigma) = -1$  otherwise.




Figure 13: A matrix and its correspondence to weighted bipartite and directed graphs. A perfect matching and its corresponding cycle cover are marked in thick red lines.

Of course,  $\det(M)$  can be computed in polynomial time (by Gaussian elimination, for example). But consider the superficially similar quantity

$$\text{perm}(M) = \sum_{\sigma \in S_n} \prod_{i=1}^n M_{i, \sigma(i)},$$

which is called the permanent of  $M$ . Then we have the following result.

**Theorem 10.3** (Valiant). *Computing  $\text{perm}(M)$  for 0-1 matrices  $M$  is #P-complete.*

Thus simply removing the apparent complicating factor of  $\text{sgn}(\sigma)$  results in a dramatic *increase* in complexity! When  $M$  is a 0-1 matrix,  $\text{perm}(M)$  has the following combinatorial interpretation. Consider a bipartite graph  $G$  with  $n$  vertices on each side of the bipartite split, where there is an edge between the  $i$ 'th vertex on the left-hand side and the  $j$ 'th vertex on the right-hand side if and only if  $M_{ij} = 1$ . Then each term  $\prod_{i=1}^n M_{i, \sigma(i)}$  is equal to 1 if and only if  $\sigma$  corresponds to a perfect matching in  $G$  – i.e. a set of  $n$  edges such that each vertex is included in exactly one edge. Hence computing the permanent of a 0-1 matrix is equivalent to counting perfect matchings in a bipartite graph, which implies that the problem is in #P. As we saw near the start of the course, the decision variant of this problem (i.e. the problem of determining whether a bipartite graph  $G$  has a perfect matching) is in P, so generalising from decision to counting makes the problem considerably more difficult. More generally, we can associate any matrix  $M$  with a bipartite undirected graph  $G$ , by putting (possibly negative) weights on each edge of  $G$ ;  $\text{perm}(M)$  is then a sum over perfect matchings in  $G$ , where each matching is weighted by the product of its edge weights.

Another interpretation of the permanent, which will be important later, is in terms of *cycle covers*. A cycle cover of  $G$  is a set of cycles in  $G$  such that each vertex is contained in exactly one cycle. It is easy to convince oneself that, if  $G$  is a directed graph with adjacency matrix  $A$ ,  $\text{perm}(A)$  is precisely the number of cycle covers of  $G$ . Similarly to the bipartite graph interpretation, for any square matrix  $M$ ,  $\text{perm}(M)$  is equal to a sum over weighted cycle covers of a graph corresponding to  $M$ . Figure 13 illustrates these ideas.

In order to prove Theorem 10.3, we first prove that computing the permanent of general integer matrices  $M$  is #P-hard, and then reduce to the case where  $M$  is a 0-1 matrix.

**Theorem 10.4.** *Computing  $\text{perm}(M)$  for integer matrices  $M$  is #P-hard.*

*Proof.* We will reduce #3-SAT to computing the permanent. Given a 3-CNF boolean formula  $\phi$  containing  $m$  literals, we will produce a matrix  $M$  such that  $\text{perm}(M) = 4^m \#\phi$ , where  $\#\phi$  is the





Figure 14: Variable and clause gadgets. Undirected edges (i.e. without arrows) correspond to an arc in each direction. In each cycle cover, exactly one of the two rightwards-moving arcs of each variable gadget must be included, corresponding to true or false. Also, at least one of the external edges of each clause gadget must be *excluded*, corresponding to a variable satisfying that clause.




Figure 15: Constructing a graph whose cycle covers correspond to satisfying assignments to the formula  $(x_1 \vee \neg x_2 \vee x_3) \wedge (\neg x_1 \vee x_2 \vee x_3)$ .




Figure 16: Constructing an XOR gate. The graph  $G_{\text{XOR}}$  is intended to ensure that there is either an edge from  $u$  to  $v$ , or from  $u'$  to  $v'$ , but not both. Here we describe  $G_{\text{XOR}}$  by its graph and its adjacency matrix, where vertices are ordered  $(s, a, b, t)$ .

number of satisfying assignments of  $\phi$ . In other words, we will construct a weighted graph  $G$  such that the sum over weighted cycle covers of  $G$  is equal to  $4^m \#\phi$ .

The construction will use three kinds of gadget: variable, clause and XOR gadgets. Each gadget corresponds to a subgraph as illustrated in Figures 14 and 16. First consider the graph formed by taking one variable gadget  $G_i$  per variable  $x_i$ , and one clause gadget  $G_C$  per clause  $C$ . We would like to make each cycle cover of this graph correspond to a satisfying assignment of  $\phi$ . Label each of the three external arcs of a clause gadget with a literal used in that clause. An external arc being included in a particular cycle cover of  $G_C$  is intended to correspond to that literal *not* satisfying  $C$ . If a clause contains fewer than three variables, the remaining arcs are left unlabelled. Observe that any cycle cover of  $G_C$  must include at most two of the external arcs<sup>1</sup>. Further, it is easy to convince oneself that every choice of a proper subset of the external arcs of  $G_C$  (including the empty set) corresponds to exactly one cycle cover of  $G_C$ .

Imagine we had a (magical!) “XOR gadget” which we could apply to any pair of arcs, and which would enforce the constraint that exactly one of the two arcs is included in any cycle cover. Given such a gadget, we could calculate  $\#\phi$  as follows. Whenever a variable  $x_i$  appears in a clause  $C$ , apply the XOR gadget between the external arc of  $G_C$  corresponding to  $x_i$ , and either the “false” arc in  $G_i$  (if  $x_i$  appears negated in  $C$ ), or the “true” arc in  $G_i$  (if  $x_i$  appears un-negated in  $C$ ). Then each assignment to variables  $x_1, \dots, x_n$  would correspond to exactly one cycle cover (if  $x$  satisfies  $\phi$ ), or no cycle covers (if  $x$  does not satisfy  $\phi$ ), so the permanent of the whole graph would equal  $\#\phi$ . Figure 15 illustrates this construction.

We will simulate such an XOR gadget using the construction  $G_{\text{XOR}}$  illustrated in Figure 16, and analyse the permanent of the resulting graph. The permanent is a sum over weighted cycle covers, and can hence be split up in terms of sums over cycle covers that use different subsets of the four arcs  $u \rightarrow s$ ,  $s \rightarrow v'$ ,  $t \rightarrow v$ ,  $u' \rightarrow t$ .

First, the sum over cycle covers that use *none* of these four arcs is equal to 0. This is because the permanent of  $G_{\text{XOR}}$ 's weighted adjacency matrix is equal to 0 (as can be checked), and the permanent of a graph made of two disjoint subgraphs is equal to the product of their respective permanents. Similarly, the sum over cycle covers that use either of the pairs of arcs  $(u \rightarrow s, s \rightarrow v')$  and  $(u' \rightarrow t, t \rightarrow v)$ , or both of these pairs, is also equal to 0. Thus the only cycle covers which contribute to the permanent of the whole graph are those which either contain an arc  $u \rightarrow s$  and an arc  $t \rightarrow v$ , or contain an arc  $u' \rightarrow t$  and an arc  $s \rightarrow v'$ , but not both of these pairs. One can calculate that in either of these cases, the weighted sum of cycle covers gets a multiplicative contribution of 4 from  $G_{\text{XOR}}$ .

This essentially completes the proof: each cycle cover which does not satisfy all of the XOR constraints gives a contribution of 0 to the overall sum, and each cycle cover which does satisfy all the constraints gives a contribution of  $4^m$  (there are  $m$  XOR gadgets in total, each contributing a factor of 4). Therefore, the permanent of the adjacency matrix of  $G$  is equal to  $4^m \#\phi$ .  $\square$

The above proof appeared to crucially rely on some clever cancellations between positive and negative entries in  $M$ . Perhaps surprisingly, we now show that the permanent remains hard even when restricted to 0-1 matrices.

*Proof of Theorem 10.3.* First, we observe that any positive integer weights in the graph  $G$  corresponding to  $M$  can be replaced with 1's using the construction illustrated in Figure 17; even

---

<sup>1</sup>One may be uneasy that  $G_C$  and  $G_i$  contain multiple arcs between certain nodes. These extra arcs will disappear shortly; in any case, they can be subsumed into one weighted arc.




Figure 17: Simulating positive weights 2, 3 and  $2^n$  respectively.

exponentially large weights can be dealt with by attaching subgraphs in series. However, there are also some weights equal to  $-1$ . To deal with these, we change the problem under consideration: rather than computing the permanent, we instead consider computing the permanent modulo  $N$ , for some integer  $N$ . This problem is clearly no easier than computing the permanent, so if we can show that  $\#3\text{-SAT}$  reduces to computing the permanent modulo  $N$ , we have shown that computing the permanent itself is  $\#\text{P}$ -hard. When working modulo  $N$ ,  $-1 \equiv N - 1$ . Take  $N$  to be a large integer,  $N = 2^n + 1$ , where we pick  $n$  such that the permanent of  $M$  does not exceed  $N$  ( $n = 8m$  works, for example). Now the  $-1$ 's in  $M$  can be replaced with  $2^n$ 's, which can be simulated as before. The number of satisfying assignments of the original formula is then precisely equal to the permanent of the resulting matrix, modulo  $N$ .  $\square$

## 10.2 Historical notes and further reading

The complexity class  $\#\text{P}$  was introduced by Leslie Valiant in 1979, in the same paper in which he proved that computing the permanent is  $\#\text{P}$ -complete. He received the 2010 Turing Award, the highest award in computer science, partly for this work (but also for much more!). Arora-Barak (chapter 17) and Papadimitriou (chapter 18) have nice discussions on  $\#\text{P}$ -hardness and counting problems. In particular, the proof given here that the permanent is  $\#\text{P}$ -complete is based on Papadimitriou's.

Valiant's result shows that counting perfect matchings in bipartite graphs is  $\#\text{P}$ -hard. Remarkably, it turns out to be possible to count perfect matchings in arbitrary graphs in polynomial time – as long as the graph is planar (i.e. can be drawn in the plane with no edges crossing). The result that this is possible is known as the FKT algorithm<sup>1</sup> and has significant applications in statistical physics.

---

<sup>1</sup>See <http://www.damtp.cam.ac.uk/user/am994/presentations/matchings.pdf>, for example, for a discussion of this beautiful algorithm.




Figure 18: Some important complexity classes.

## 11 Complexity classes: the summary

Figure 18 summarises known relationships between many of the complexity classes we have met so far. All of these inclusions (with the possible exception of  $P \subseteq BPP$ ) are believed to be strict, but there is no proof that this is the case. Observe that we have the chain of inclusions

$$L \subseteq NL \subseteq P \subseteq NP \subseteq PSPACE.$$

By the Space Hierarchy Theorem, we know that  $L \neq PSPACE$ . Therefore, at least one of these inclusions must be strict – but we do not know which one.




Figure 19: A circuit  $C$  computing a boolean function of 3 variables. One can verify that in fact  $C(x) = (x_1 \vee x_2) \wedge x_3$ .

## 12 Circuit complexity

Turing machines are an (arguably!) nice model for reasoning about computation in a very general sense. However, modern computers do not usually look much like Turing machines, but instead are built from circuits of electronic components. In order to model this more directly, we consider the circuit model of computation.

For  $n \in \mathbb{N}$ , an  $n$ -input, single-output boolean circuit  $C$  is a directed acyclic graph which has one sink (node with out-degree 0). Sources (i.e. nodes with in-degree 0) are labelled with variables  $x_1, \dots, x_n \in \{0, 1\}$ , or with one of the constants 0 or 1, and are called input nodes; the sink is called the output node. Non-sources are called gates and labelled with one of  $\wedge$ ,  $\vee$  or  $\neg$  (i.e. AND, OR and NOT). Nodes labelled with  $\wedge$  or  $\vee$  have in-degree (also called *fan-in*) 2, and nodes labelled with  $\neg$  have fan-in 1. Similarly, the out-degree of a node is called its fan-out. The size of  $C$  is the number of nodes in  $C$ , and is written  $|C|$ .

$C$  computes a boolean function  $C : \{0, 1\}^n \rightarrow \{0, 1\}$  in the natural way. Each node has a value, defined recursively as follows: the value of an input node is (depending on its label) the value of its corresponding variable  $x_i$ , or the constant 0 or 1, and the value of other nodes is determined by applying the logical operation labelling that node to the values of the nodes connected to it.  $C(x)$  is then defined to be the value of the output node on input  $x$ . See Figure 19 for an example of a circuit.

Observe that every boolean formula can be written as a boolean circuit. The universality of boolean formulae (Theorem 5.1) thus implies that boolean circuits can compute any boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ . However, the construction of Theorem 5.1 may use as many as  $\Omega(n2^n)$  gates (there can be  $\Omega(2^n)$  clauses, and each clause uses  $n$  gates). This may appear somewhat inefficient; ideally, we would like circuits of size polynomial in  $n$ . However, this is not possible for every function  $f$ .

**Theorem 12.1.** *For every integer  $n > 1$ , there exists a boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  which cannot be computed by any circuit of size at most  $2^n/(2n)$ .*

*Proof.* The proof is a counting argument. There are  $2^{2^n}$  boolean functions on  $n$  bits. We will upper bound the number  $C_{n,m}$  of circuits on  $n$  bits with  $m$  gates. Each such circuit is defined by choosing which type of gate each node computes, and which gates feed into each node. For each gate we have at most  $(n+5)m^2$  possibilities for these, so we have at most  $((n+5)m^2)^m$  different circuits in total. This is a rough over-estimate as many of these possibilities will not lead to valid circuits.

We now simply observe that, for  $m \leq 2^n/(2n)$ ,

$$\log_2 C_{n,m} \leq \frac{2^n}{2n} \log_2 \left( \left( \frac{n+5}{4n} \right) 2^{2n} \right) < 2^n = \log_2 2^{2^n},$$

valid for  $n \geq 2$ . □

Observe that the same proof technique even shows that all but an exponentially small fraction of functions  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  require exponential-size circuits. Nevertheless, no explicit example of a family of such functions is known!

## 12.1 Polynomial-size circuits

We now connect the circuit model with our previously studied notion of complexity classes.

Let  $T : \mathbb{N} \rightarrow \mathbb{N}$  be a function. A  $T(n)$ -size circuit family  $\mathcal{C}$  is a sequence  $C_0, C_1, \dots$  of boolean circuits, where  $C_n$  has  $n$  inputs and one output, and  $|C_n| \leq T(n)$  for all  $n$ . Let  $\mathcal{L} \subseteq \{0, 1\}^*$  be a language. We say that a circuit family  $\mathcal{C}$  decides  $\mathcal{L}$  if, for every  $n$  and every  $x \in \{0, 1\}^n$ ,  $x \in \mathcal{L} \Leftrightarrow C_n(x) = 1$ . We further say that  $\mathcal{L} \in \text{SIZE}(T(n))$  if there exists a  $T(n)$ -size circuit family  $\mathcal{C}$  such that  $\mathcal{C}$  decides  $\mathcal{L}$ .

Observe that, by contrast with the Turing machine model, we do not demand that a single circuit works for arbitrary input sizes  $n$ . Here, we are allowed to have a completely different circuit for each  $n$ . By analogy with the Turing machine model, we will be interested in “small” classes of circuits. Let P/poly be the class of languages decided by polynomial-sized circuit families, i.e.

$$\text{P/poly} = \bigcup_{c>0} \text{SIZE}(n^c).$$

We first show that polynomial-sized circuits are at least as powerful as polynomial-time Turing machines.

**Theorem 12.2.**  $\text{P} \subseteq \text{P/poly}$ .

*Proof.* Consider a language  $\mathcal{L} \in \text{P}$  and let  $x \in \{0, 1\}^n$ . We will show that there is a circuit  $C$  of size  $\text{poly}(n)$  which computes the function  $C(x) = 1 \Leftrightarrow x \in \mathcal{L}$ . The proof is similar to the proof of the Cook-Levin theorem.

Let  $M$  be a Turing machine that decides  $\mathcal{L}$  in time  $n^k$ , for some constant  $k$ , and let  $T$  be the tableau (defined as in the Cook-Levin theorem) describing the computational steps  $M$  takes on input  $x$ . The first row and the first and last columns of  $T$  are known in advance (being determined only by  $x$ , equal to a column of  $\triangleright$ 's, and equal to a column of  $\square$ 's, respectively). Because of the locality of Turing machines, each other cell  $T_{ij}$  depends only on three cells ( $T_{i-1,j-1}$ ,  $T_{i-1,j}$  and  $T_{i+1,j-1}$ ) in the row above. Thus the contents of each cell  $T_{ij}$  is a function of three other cells. Let  $S$  be the set of symbols used in the tableau. Each symbol can be encoded as a fixed-length binary string (of length  $m$ , say), so the tableau can be thought of as an  $n^k \times n^k \times m$  array  $T_{ij\ell}$ . By the above observation, each entry  $T_{ij\ell}$  depends only on  $3m$  other entries. This means that we can write

$$T_{ij\ell} = f_\ell(T_{i-1,j-1,1}, \dots, T_{i-1,j-1,m}, T_{i,j-1,1}, \dots, T_{i+1,j-1,m}),$$

where  $f_\ell : \{0, 1\}^{3m} \rightarrow \{0, 1\}$ . Each  $f_\ell$  can thus be described by a circuit  $C_\ell$ , which is of constant size and depends only on  $M$ . The overall circuit  $C$  consists of the repeated application of these

individual circuits  $C_\ell$ , with one application for each cell  $(i, j)$ . For each entry  $T_{ij}$ , the input gates of  $C_\ell$  are attached to the output gates of  $T_{i-1, j-1}$ ,  $T_{i-1, j}$  and  $T_{i+1, j-1}$ . Finally, the output of  $C$  is determined by making  $C$  output 1 if and only if the last row of the tableau corresponds to  $M$  outputting 1.

It is clear that the resulting circuit is of size  $\text{poly}(n)$ . Further note that (as will be important later) the above reduction can be performed in space  $O(\log n)$ . Each circuit  $C_\ell$  is fixed (i.e. depends only on  $M$ ). Describing  $C$  thus requires writing down the input gates (which depend only on  $|x|$ ), generating many copies of  $C_\ell$  and identifying correct inputs and outputs of these circuits, and writing down the output gate of  $C$ . Because of the regular structure of  $C$ , this can be done in logarithmic space.  $\square$

The above inclusion is in fact proper in a very strong sense: there even exist languages in  $\text{P/poly}$  which are undecidable! Let UHALT be the following language.

UHALT =  $\{1^n : \text{the binary expansion of } n \text{ encodes a pair } (M, x) \text{ such that } M \text{ halts on input } x\}$ .

It is clear that UHALT is undecidable. On the other hand, we have the following result.

**Theorem 12.3.** *UHALT is in P/poly.*

*Proof.* The theorem will follow from showing that every *unary language* is in  $\text{P/poly}$ , where a unary language  $\mathcal{L} \subseteq \{0, 1\}^*$  is a language of the form  $\mathcal{L} \subseteq \{1^n : n \in \mathbb{N}\}$ . We describe a circuit family  $\mathcal{C}$  for deciding  $\mathcal{L}$ , where each member  $C_n$  of the family behaves as follows. If  $1^n \in \mathcal{L}$ , then the circuit computes the AND of the  $n$  input bits (which can be done with a linear size circuit). Otherwise, the circuit always outputs 0.  $\square$

In order to conclude that UHALT  $\in \text{P/poly}$ , it was sufficient to show that there exists a circuit to solve any given instance of UHALT, without needing to actually construct such a circuit. We would like to avoid such pathological cases in our model of circuit complexity. To do so, we introduce the notion of uniformly generated circuits.

A circuit family  $\mathcal{C} = C_0, C_1, \dots$  is said to be  $L$ -uniform if there is a logarithmic-space Turing machine which, on input  $1^n$ , outputs a description of the circuit  $C_n$  (i.e. if there exists an implicitly logspace-computable function which maps  $1^n$  to  $C_n$ ). This type of circuit turns out to be precisely equivalent to  $\text{P}$ .

**Theorem 12.4.** *A language  $\mathcal{L} \subseteq \{0, 1\}^*$  can be decided by an  $L$ -uniform circuit family if and only if  $\mathcal{L} \in \text{P}$ .*

*Proof.* ( $\Rightarrow$ ): If  $\mathcal{L} \subseteq \{0, 1\}^*$  is decided by an  $L$ -uniform circuit family  $\mathcal{C}$ , then for any  $x \in \{0, 1\}^*$ , we can determine whether  $x \in \mathcal{L}$  using a polynomial-time Turing machine by first constructing the required circuit  $C_{|x|}$ , then evaluating it.

( $\Leftarrow$ ): Follows from the proof of Theorem 12.2, which gives a logarithmic-space construction of a circuit  $C_{|x|}$  which outputs 1 if and only if  $x \in \mathcal{L}$ .  $\square$

The same idea allows us to give a natural complete problem for  $\text{P}$ . As with the class  $\text{NL}$  (Section 8.2), to formulate a sensible notion of complete problems for  $\text{P}$  it is necessary to use a weaker notion of reductions than for  $\text{NP}$  (which motivates the above introduction of the idea of  $L$ -uniformity). As with  $\text{NL}$ , the notion of reductions we use is implicit logarithmic space reductions. We say that

a language  $\mathcal{A}$  is P-complete if  $\mathcal{A} \in \text{P}$ , and for all  $\mathcal{B} \in \text{P}$ ,  $\mathcal{B} \leq_L \mathcal{A}$ . P-complete problems are the “hardest” problems in P; indeed, if  $\mathcal{L}$  is P-complete and  $\mathcal{L} \in \text{L}$ , then  $\text{P} = \text{L}$ .

The CIRCUIT VALUE problem is defined as follows. The input is a pair  $(C, x)$ , where  $C$  is the description of an  $n$ -input circuit and  $x \in \{0, 1\}^n$ . The problem is to decide whether  $C(x) = 1$ .

**Theorem 12.5.** CIRCUIT VALUE is P-complete.

*Proof.* Evaluating  $C(x)$  can clearly be done in polynomial time; for the other direction, Theorem 12.2 gives an implicitly logspace-computable reduction from any problem in P to CIRCUIT VALUE.  $\square$

Similarly, we can give a characterisation of the class NP in terms of circuits. The CIRCUIT SAT problem is defined as follows: the input is a description of an  $n$ -input circuit  $C$ , and the problem is to decide whether there exists  $x \in \{0, 1\}^n$  such that  $C(x) = 1$ . As boolean formulae are special cases of circuits, CIRCUIT SAT is at least as hard as 3-SAT. We now give an alternative proof of this fact which follows directly from the definition of NP.

**Theorem 12.6.** CIRCUIT SAT is NP-complete.

*Proof.* CIRCUIT SAT is clearly in NP, as given  $x \in \{0, 1\}^n$ , we can evaluate  $C(x)$  in polynomial time. For the other direction, if  $\mathcal{L} \in \text{NP}$  then there is a polynomial-time Turing machine  $M$  and a polynomial  $p$  such that  $x \in \mathcal{L}$  if and only if  $M(x, w) = 1$  for some  $w \in \{0, 1\}^{p(n)}$ . The proof of Theorem 12.2 shows that, given the pair  $(\underline{M}, x)$ , we can write down a circuit  $C_x$  in polynomial time such that  $C_x(w) = M(x, w)$  for all  $w \in \{0, 1\}^{p(n)}$ . Thus solving CIRCUIT SAT for  $C_x$  allows us to determine whether  $x \in \mathcal{L}$ .  $\square$

This now allows us to go full circle and prove NP-hardness of 3-SAT using CIRCUIT SAT – i.e. to give an alternative proof of the Cook-Levin theorem.

**Theorem 12.7.** CIRCUIT SAT  $\leq_P$  3-SAT.

*Proof sketch.* Given a circuit  $C$  containing AND, OR and NOT gates, we produce a 3-CNF formula  $\phi$  as follows. Each node  $v_i$  has a corresponding variable  $z_i$  in  $\phi$ . We then add clauses to  $\phi$  such that the input and output constraints of each gate are satisfied. For example, if  $v_i = v_j \wedge v_k$  is an AND gate, we want to add the condition  $(z_i = z_j \wedge z_k)$ , which can easily be written as a CNF formula; the same applies for OR and NOT. Finally, if the output node is  $v_o$ , we add a clause  $(z_o)$  to specify that the output should be 1. It is easy to see that  $\phi$  is satisfiable if and only if  $C$  is satisfiable.  $\square$

As the reductions occurring in these two theorems are parsimonious, this finally proves the claim made in Section 10 that the counting problem #3-SAT is #P-complete.

## 12.2 Restricted depth circuits

We have seen that size in the model of uniformly generated circuits is roughly equivalent to time in the world of Turing machines. However, the circuit model allows us to consider other measures of complexity which do not necessarily correspond to natural parameters in the Turing machine model. A particularly interesting such measure is circuit depth, where the depth of a circuit  $C$  is the maximal number of nodes on a path from an input node to the output node. We now define two classes of depth-limited computations.






Figure 20: Recursive construction for computing PARITY, and this construction mapped onto the standard circuit model.

- For every  $d$ , a language  $\mathcal{L} \subseteq \{0, 1\}^*$  is in the class  $\text{NC}^d$  if  $\mathcal{L}$  can be decided by an L-uniform family of circuits  $\{C_n\}$  such that each  $C_n$  has size  $\text{poly}(n)$  and depth  $O(\log^d n)$ . We write  $\text{NC} = \bigcup_{d \geq 0} \text{NC}^d$ .
- For every  $d$ , the class  $\text{AC}^d$  is defined similarly, but with the extension that we allow gates in each circuit  $C_n$  to have unbounded fan-in. That is, each OR and AND gate can be applied to arbitrarily many input bits. We write  $\text{AC} = \bigcup_{d \geq 0} \text{AC}^d$ .

Thus, for each  $d$ , the class  $\text{AC}^d$  is at least as powerful as  $\text{NC}^d$ . Indeed, we have

$$\text{NC}^d \subseteq \text{AC}^d \subseteq \text{NC}^{d+1}$$

(you will show the second inclusion in an exercise). The class  $\text{NC}^0$  is very limited as the output of any circuit in this class can depend only on a constant number of input bits;  $\text{AC}^0$  does not have this restriction and is arguably the simplest interesting class of circuits which one could consider.

There are two main motivations for studying such classes. One motivation is their connection to reality; the class  $\text{NC}$  in particular corresponds quite well to a model of massively parallel computing where we have access to a large number of processors with efficient communication between them (such a model was not realistic for some time after  $\text{NC}$  was first studied, but nowadays is not too far from the truth). Another motivation is the fact that it seems difficult to prove a super-polynomial lower bound on circuit size for an explicit family of functions. Restricting ourselves to considering circuits in  $\text{NC}$  or  $\text{AC}$  might allow better lower bounds to be proven.

In particular, we will consider the problem PARITY, which is defined as follows: given  $x \in \{0, 1\}^n$ , does  $x$  contain an odd number of 1's? Observe that this problem is in  $\text{NC}^1$ : given an  $n$ -bit string  $x$ , solving PARITY is the same as computing  $x_1 \oplus x_2 \oplus \dots \oplus x_n$ , which can be done in depth  $O(\log n)$  by a recursive construction illustrated in Figure 20.

However, we have the following result.

**Theorem 12.8.** PARITY  $\notin \text{AC}^0$ . Hence  $\text{AC}^0 \neq \text{NC}^1$ .

The class  $\text{AC}^0$  appears very weak, so it is not clear that this is an interesting result. However, it is almost the strongest lower bound known on circuit size!

The basic idea behind Theorem 12.8 is as follows. We will show that any function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  computed by a constant-depth, polynomial-size circuit has the property that we can force it to become constant by fixing strictly less than  $n$  of its input variables. As PARITY does not

have this property,  $\text{PARITY} \notin \text{AC}^0$ . To gain some intuition, consider the special case where  $f$  is computed by a  $k$ -CNF formula for  $k < n$  (recall that a  $k$ -CNF (resp.  $k$ -DNF) formula is a boolean formula which is an AND of ORs (resp. OR of ANDs) where each OR (resp. AND) involves at most  $k$  variables). We can make such a formula constant by forcing one clause to evaluate to 0 or 1, respectively, which requires fixing at most  $k$  of the variables. Thus  $\text{PARITY}$  cannot be computed by such a formula.

The proof of Theorem 12.8 will be based on extending this idea to general circuits, using the following “switching lemma”, which we state without proof. The lemma uses the concept of random restrictions of a boolean formula. If  $f$  is a boolean function and  $\rho$  is a partial assignment (known as a restriction) to the variables of  $f$ , let  $f|_\rho$  denote the function obtained by fixing the variables of  $\rho$ .

**Lemma 12.9** (Håstad’s Switching Lemma). *Let  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  be expressible as a  $k$ -DNF, and let  $f|_\rho$  be a random restriction that assigns random values to  $t$  randomly selected input bits. Then, for every  $s \geq 2$ ,*

$$\Pr_\rho[f|_\rho \text{ is not expressible as } s\text{-CNF}] \leq \left((1 - t/n)k^{10}\right)^{s/2}.$$

By symmetry (i.e. applying Lemma 12.9 to  $\neg f$ ), we get the same result with the terms DNF and CNF interchanged.

*Proof of Theorem 12.8 (sketch).* We assume that we are given an  $\text{AC}^0$  circuit  $C$  of depth  $d$  in the following form (any circuit can be converted into such a form).

- The circuit is a tree (i.e. all gates have fan-out 1);
- All  $\neg$  gates are at the input of the circuit;
- At each level of the tree, all gates are either  $\vee$  or  $\wedge$ ;
- The first level after the input consists of  $\vee$  gates of fan-in 1.

Let  $n^c$  be an upper bound on the number of gates in  $C$ . We randomly restrict more and more variables; the idea is that each restriction reduces the depth by 1 while retaining constant fan-in at the first level of the circuit. Let  $n_i$  be the number of unrestricted variables after step  $i$ . At step  $i + 1$  we randomly restrict  $\lceil n_i - \sqrt{n_i} \rceil$  variables. As  $n_0 = n$ ,  $n_i \approx n^{2^{-i}}$ . Set  $k_i = 10c2^i$ . We aim to show by induction that, with high probability, after restriction  $i$  we have a circuit of depth at most  $(d - i)$  and with fan-in at most  $k_i$  at the first level.

Suppose the first level contains  $\vee$  gates, so the second level contains  $\wedge$  gates. Assuming that the function computed by each  $\wedge$  gate is a  $k_i$ -CNF, by Lemma 12.9, after the  $(i + 1)$ ’th step it will be expressible as a  $k_{i+1}$ -DNF with probability at least

$$1 - \left(k_i^{10} n^{-2^{-i-1}}\right)^{k_{i+1}/2},$$

which is at least  $1 - (10n^c)^{-1}$  for large enough  $n$  (observing that  $k_i$  is constant with respect to  $n$ ). As the bottom gate of a DNF formula is  $\vee$ , we can merge the gates at the second level with the level below, decreasing the circuit’s depth by 1. On the other hand, if the first level contained  $\wedge$  gates, we can transform the  $k_i$ -DNF of the next level into a  $k_{i+1}$ -CNF using the same idea.

We repeat this process  $d - 2$  times. With constant probability we end up with a depth 2 circuit with fan-in at most  $k_{d-2}$  at the top level. To see this, note that we apply Lemma 12.9 once per gate, so using a union bound the probability that all of the applications were successful is at least  $1 - n^c(10n^c)^{-1} = 9/10$ . A depth 2 circuit is just a  $k_{d-2}$ -CNF or  $k_{d-2}$ -DNF formula. But if we fix at most  $k_{d-2} = O(1)$  of the variables in such a formula, we can make the formula constant (by fixing one clause to 0 or 1, respectively). The PARITY function is not constant under any restriction of fewer than  $n$  variables, so PARITY  $\notin AC^0$ .  $\square$

### 12.3 Circuits and randomised algorithms

Recall that it is not known whether  $BPP = P$ . However, in the circuit model one can indeed prove a result of this nature.

**Theorem 12.10.**  $BPP \subset P/\text{poly}$ .

*Proof.* Let  $\mathcal{L} \in BPP$ . From the proof of Lemma 9.8, this implies that there exists a polynomial-time Turing machine  $M$  which, on inputs of size  $n$ , uses some random bits  $r \in \{0, 1\}^m$ , for some  $m \geq 0$ , and for all  $x \in \{0, 1\}^n$ ,  $\Pr_r[M(x, r) \neq f_{\mathcal{L}}(x)] \leq 2^{-n-1}$ . (Recall that  $f_{\mathcal{L}}(x) = 1$  if  $x \in \mathcal{L}$ , and  $f_{\mathcal{L}}(x) = 0$  if  $x \notin \mathcal{L}$ .) Call a string  $r \in \{0, 1\}^m$  bad for  $x$  if  $M(x, r) \neq f_{\mathcal{L}}(x)$ , and otherwise call  $r$  good for  $x$ . For each  $x$ , at most  $2^{m-n-1}$  strings  $r$  are bad for  $x$ . Thus, summing over all  $x$ , there are at most  $2^{m-1}$  strings which are bad for *some*  $x$ . In particular, for each  $n$  there is a string  $g \in \{0, 1\}^m$  which is good for every  $x \in \{0, 1\}^n$ . We can fix this string  $g$  as input to  $M$  and hence obtain a circuit  $C_n$  whose size is polynomial in the time used by  $M$ , and such that  $C(x) = f_{\mathcal{L}}(x)$  for all  $x \in \{0, 1\}^n$ .  $\square$

Of course, it cannot be the case that in fact  $BPP = P/\text{poly}$ , as  $P/\text{poly}$  contains undecidable problems. Theorem 12.10 can be seen as evidence that  $BPP = P$ . Additional evidence is provided by the following result, which we state informally: if there exists a language  $\mathcal{L} \in DTIME(2^{O(n)})$  such that deciding  $\mathcal{L}$  requires exponentially large circuits, then  $BPP = P$ . See Arora-Barak chapter 20 for a discussion and proof of this claim.

### 12.4 Natural proofs and P vs. NP

As  $P \subseteq P/\text{poly}$ , one plausible way to prove  $P \neq NP$  is to prove  $NP \not\subseteq P/\text{poly}$  (again, we cannot have  $P/\text{poly} = NP$  because  $P/\text{poly}$  contains undecidable problems). This could be a hopeful strategy because circuits may seem easier (or more “concrete”) to reason about than Turing machines. The most direct way to prove this separation would be to show that some problem believed not to be in  $P$  (such as the INTEGER FACTORISATION problem, or any NP-complete problem) requires circuits of super-polynomial size. Indeed, the result that PARITY  $\notin AC^0$  seems to represent progress in this direction.

Ideally, we could imagine that we would develop a test which we could apply to the truth table of any boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , which would, in time polynomial in the size of the truth table, output either “easy” ( $f$  has a circuit of size at most  $n^k$ , for some  $k$ ) or “hard” ( $f$  does not have such a circuit). If we could then find a family of functions  $f$  known to be in NP, but for which this test output “hard”, we would have shown  $NP \not\subseteq P/\text{poly}$ . As well as helping to resolve complexity-theoretic issues, such a test would be very useful in practice.

Unfortunately, it has been shown that (modulo some reasonable assumptions) such a test cannot exist, significantly restricting the potential for this approach to prove  $P \neq NP$ . The proof of this

result is of some conceptual interest in itself as it links two apparently (?) disparate areas: hardness and pseudorandomness.

In order to state the result more formally, we will need some definitions. Let  $F_n$  be the set of boolean functions  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ . A *property*  $P_n$  is a function  $P_n : F_n \rightarrow \{0, 1\}$ . We say that  $P_n$  is *natural* if it satisfies the following two properties.

- **(Constructivity)** For any  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , computing  $P_n(f)$  can be done in time polynomial in the size of the truth table of  $f$ , i.e. in time  $2^{O(n)}$ .
- **(Largeness)** At least a  $2^{-O(n)}$  fraction of functions  $f \in F_n$  satisfy  $P_n(f) = 1$ .

A good example of a natural property is the property “ $f$  does not become constant after fixing all but  $n^\epsilon$  variables of the input, for some constant  $\epsilon$ ”, which was used in the proof that  $\text{PARITY} \notin \text{AC}^0$ . It can be verified that this property is satisfied by random functions with high probability, and can be checked in time  $2^{O(n)}$ .

We say that  $P_n$  is *useful* against P/poly if the circuit size of any sequence of functions  $f_1, f_2, \dots$ , where  $f_n \in P_n$ , is super-polynomial in  $n$ ; i.e. for any constant  $k$ , for sufficiently large  $n$ , the circuit size of  $f_n$  is larger than  $n^k$ . If a natural property  $P_n$  exists which is useful against P/poly, we therefore have an algorithm which we can apply to boolean functions to certify that they do not have small circuits. Equivalent definitions can be formulated for proofs useful against other circuit complexity classes, but we will only consider P/poly here.

We will also need the concept of one-way functions, which are functions which are easy to compute but hard to invert. The field of cryptography is essentially based around the idea of one-way functions (although we do not know if any such functions exist!).

A polynomial-time computable function  $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$  is said to be a *subexponentially strong one-way function* if there is an  $\epsilon > 0$  such that, for every probabilistic algorithm  $\mathcal{A}$  running in time  $2^{n^\epsilon}$ , there is a function  $\delta : \mathbb{N} \rightarrow [0, 1]$  such that, for every  $n$ ,

$$\Pr_{x \in \{0, 1\}^n} [\mathcal{A}(f(x)) = y \text{ for some } y \text{ such that } f(y) = f(x)] < \delta(n),$$

and  $\delta(n) < n^{-c}$  for every constant  $c$  and sufficiently large  $n$ . In other words, any probabilistic algorithm running in time  $2^{n^\epsilon}$  has super-polynomially small probability of finding some  $y$  such that  $f(y) = f(x)$ . We think of  $f$  as encoding its input  $x$  in such a way that it is hard for any adversary to decode it. It is not obvious that one-way functions should exist; indeed, if they do exist, then  $\text{P} \neq \text{NP}$ . However, a number of natural functions are conjectured to be one-way. A simple example is integer multiplication: inverting the function  $f$  which treats its input as two  $n$ -bit numbers and outputs their product is essentially the INTEGER FACTORISATION problem. (Recall that the best known algorithm for this problem uses time  $2^{O(n^{1/3}(\log n)^{2/3})}$ .)

We are finally ready to state the main result about natural proofs.

**Theorem 12.11.** *If there exists a natural property  $P_n$  which is useful against P/poly, subexponentially strong one-way functions cannot exist.*

We prove this theorem by appealing to a result connecting pseudo-random function generators and one-way functions, which we will state without proof.

A *pseudo-random function generator* is a polynomial-time computable function  $f : \{0, 1\}^{n^c} \rightarrow F_n$ , for some constant  $c > 1$ .  $f$  takes as input an  $n^c$ -bit “key”  $s$ , and outputs a function  $f(s) \in F_n$

which is indistinguishable from a random function to any probabilistic algorithm  $\mathcal{A}$  that runs in time polynomial in the truth table of  $f(s)$ , i.e. in time  $2^{O(n)}$ . Formally,

$$\left| \Pr_{g \in F_n} [\mathcal{A}(g) = 1] - \Pr_{s \in \{0,1\}^{n^c}} [\mathcal{A}(f(s)) = 1] \right| \leq 2^{-O(n)}.$$

In particular, observe that  $\mathcal{A}$  cannot just attempt to guess the key  $s$ , as this would require time  $2^{\Theta(n^c)}$ . The following lemma, which we will not prove here, states that pseudo-random function generators can be constructed from one-way functions.

**Lemma 12.12.** *Given that a family of subexponentially strong one-way functions exists, then a family of pseudo-random function generators exists.*

Based on this lemma, we can show that the existence of useful natural proofs implies the nonexistence of pseudo-random function generators, thus proving Theorem 12.11.

**Lemma 12.13.** *If there exists a natural property  $P_n$  which is useful against P/poly, pseudo-random function generators cannot exist.*

*Proof.* Assume that such a property  $P_n$  does exist. Then we can distinguish the output of such pseudo-random function generators from true randomness, i.e.:

$$\left| \Pr_{g \in F_n} [P_n(g) = 1] - \Pr_{s \in \{0,1\}^{n^c}} [P_n(f(s)) = 1] \right| \geq 2^{-O(n)}.$$

To see this, note that  $f(s)$  is polynomial-time computable, and thus cannot satisfy the property  $P_n$ . However, by the largeness property, at least a  $2^{-O(n)}$  fraction of functions will be in  $P_n$ , and hence we are able to distinguish the two cases. Note that we have used all the attributes of a natural property useful against P/poly: constructivity, largeness and usefulness.  $\square$

## 12.5 Historical notes and further reading

The proof that most boolean functions require exponential size goes back to Shannon in 1949.

NC stands for “Nick’s Class” in honour of Nick Pippenger, who first defined the class, while the A in AC stands for “alternations”. The discussion of the circuit model given here is largely based on Arora-Barak (chapter 6) and Papadimitriou (chapters 11 and 15). Arora-Barak chapter 6 includes a proof of Håstad’s Switching Lemma.

The notion of “natural proof” was invented by Razborov and Rudich in 1994, and has turned out to be a remarkably wide-ranging conceptual roadblock to proving  $P \neq NP$ . See Arora-Barak chapter 23 for a discussion. There is a vast literature concerning pseudorandomness and one-way functions. For details of the results sketched out here, a good source is the two-volume work “Foundations of Cryptography” by Oded Goldreich<sup>1</sup>.

---

<sup>1</sup><http://www.wisdom.weizmann.ac.il/~oded/foc-drafts.html>




Figure 21: A decision tree of depth 3.

### 13 Decision trees

We have seen that there are many conjectured separations between complexity classes which are not known to hold. We now change tack and consider a much simpler model than the Turing machine, where it is much easier to find provable lower bounds on the complexity of problems. Nevertheless, even in this simple model there are still many open questions.

Imagine we want to compute some (known) boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  of some (unknown) input  $x \in \{0, 1\}^n$ . Although  $x$  is unknown, we are given access to individual bits of  $x$  via an oracle  $O_x$  which, on input of an index  $i \in \{1, \dots, n\}$ , replies with the bit  $x_i$ . Our goal is to compute  $f(x)$  using the minimal number of queries to  $O_x$  in the worst case, i.e. to minimise the maximum (over  $x$ ) possible number of queries used. It is clear that any function  $f(x)$  can be computed using  $n$  queries, as this is enough to learn  $x$  completely. However, certain functions can be computed with far fewer queries.

For example, consider the function  $\text{ADDR} : \{0, 1\}^{m+2^m} \rightarrow \{0, 1\}$  where we divide the input of size  $n = m + 2^m$  into blocks of  $m$  bits and  $2^m$  bits (an “address register” and a “data register”), and define  $\text{ADDR}(i, x) = x_i$ . That is,  $\text{ADDR}$  returns the bit of the data register given by the address register. It is easy to see that, if we query every bit of the address register, this suffices to compute  $\text{ADDR}$  on every possible input and uses only  $O(\log n)$  queries.

In general, the operation of a deterministic algorithm in the query model is given by a *decision tree* (see Figure 21 for a simple example). At each step of the algorithm, we query a bit  $x_i$ , where the choice of  $i$  depends on the responses to all previous queries we have made. At the end of the algorithm, we either output 0 or 1. Thus we can express the operation of the algorithm by a binary tree whose vertices are labelled with variables  $x_i$ ,  $1 \leq i \leq n$ , and whose leaves are labelled with 0 or 1. The worst-case number of queries made by the algorithm is the depth of the tree, i.e. the maximal number of vertices labelled by variables encountered on any path from the root to a leaf. The minimal depth over all decision trees computing  $f$  is called the *decision tree complexity* of  $f$  and denoted  $D(f)$ . More concisely,

$$D(f) = \min_{T, T \text{ computes } f} \max_{x \in \{0, 1\}^n} [\text{number of variables queried by } T \text{ on input } x].$$

A function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  is called *evasive* if  $D(f) = n$ .

We have the following observations about decision tree complexity.

**Lemma 13.1.** *For any  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  that depends on all  $n$  variables,  $D(f) = \Omega(\log n)$ .*

*Proof.* Any decision tree  $T$  with depth at most  $d$  contains at most  $2^{d+1} - 1$  internal vertices, so the function computed by  $T$  can depend on at most  $2^{d+1} - 1$  variables.  $\square$

Recall that the Hamming weight of a bit-string  $y \in \{0, 1\}^n$ , which we write  $\text{wt}(y)$ , is the number of 1's in  $y$ , i.e.  $\text{wt}(y) = |\{i : y_i = 1\}|$ .

**Lemma 13.2.** *If  $|\{x : f(x) = 1\}|$  is odd, then  $f$  is evasive. Further, if  $f$  is not evasive, then exactly half the bit-strings  $x$  such that  $f(x) = 1$  have even Hamming weight.*

*Proof.* For the first part, we prove the contrapositive. For any leaf  $\ell$  at depth  $d$  in a decision tree  $T$  computing a function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , exactly  $2^{n-d}$  possible input strings  $x$  lead to  $\ell$ . If  $d < n$ , an even number of inputs  $x$  lead to  $\ell$ . As each string leads to exactly one leaf, if the depth of  $T$  is less than  $n$ ,  $|\{x : f(x) = 1\}|$  is even. For the second part, note that for each such leaf  $\ell$ , exactly half of the strings that lead to  $\ell$  have even Hamming weight.  $\square$

One method of lower bounding decision tree complexity is the adversary approach. Here we imagine that the input  $x$  is chosen bit by bit, by an adversary, in order to frustrate any possible decision tree making a small number of queries. For example, consider the  $\text{OR}_n$  function  $\text{OR}_n(x_1, \dots, x_n) = x_1 \vee x_2 \cdots \vee x_n$ ; we will show that  $D(\text{OR}_n) = n$ . Let  $T$  be a decision tree computing  $\text{OR}_n$ , and assume (without loss of generality) that  $T$  never queries the same variable twice. Whichever variable is queried by  $T$ , the adversary will always respond with 0. Even after  $n - 1$  queries, the algorithm therefore does not know if the answer should be 0 or 1.

The following technique, which is known as the method of polynomials, provides a powerful alternative tool to lower bound decision tree complexity. Any function  $f : \{0, 1\}^n \rightarrow \mathbb{R}$  can be written as a multilinear polynomial  $p : \mathbb{R}^n \rightarrow \mathbb{R}$  in  $n$  variables  $x_1, \dots, x_n \in \mathbb{R}$ , and this polynomial is unique (you will prove these claims in an exercise). Let  $\deg(f)$  be the degree of this polynomial (i.e. the largest total number of variables in any term of  $p$ ). For example:

- The function computed by the decision tree illustrated in Figure 21 has polynomial representation  $1 - x_1 - x_2 + x_1x_2 + x_2x_3 + x_2x_5 - x_1x_2x_3 - x_2x_3x_5$  and hence degree 3.
- The function  $\text{AND}_n$  defined by  $\text{AND}_n(x_1, \dots, x_n) = x_1 \wedge \cdots \wedge x_n$  has polynomial representation  $x_1x_2 \dots x_n$  and hence  $\deg(\text{AND}_n) = n$ .

**Theorem 13.3.** *For any boolean function  $f$ ,  $D(f) \geq \deg(f)$ .*

*Proof.* We show by induction that any decision tree  $T$  of depth  $d$  that computes a boolean function  $f$  gives rise to a degree  $d$  polynomial that represents  $f$ . This is clearly true for decision trees of depth 0 (which are just constant functions). So assume without loss of generality that  $x_1$  is the first variable queried by  $T$ . Then  $f(x)$  can be written as  $f(x) = (1 - x_1)f_0(x) + x_1f_1(x)$  for some functions  $f_0, f_1$  computed by decision trees of depth at most  $d - 1$ . As the degree of  $f_0$  and  $f_1$  is at most  $d - 1$ , the degree of  $f$  is at most  $d$ .  $\square$

### 13.1 Nondeterministic decision tree complexity

Just as in the model of time complexity, there is a nondeterministic analogue of the deterministic decision tree model. It is most natural to express the idea of nondeterminism using the concept of certificates, which are defined as follows.

Let  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  be a boolean function, and let  $C : S \rightarrow \{0, 1\}$  be an assignment to some subset  $S \subseteq [n]$  of the input variables. We say that  $C$  is a 0-certificate (resp. 1-certificate) if  $f(x) = 0$  (resp.  $f(x) = 1$ ) whenever the variables in  $S$  are assigned according to  $C$ . We further say that  $C$  is a 0-certificate for  $f$  on  $x$ , for some  $x \in \{0, 1\}^n$  such that  $f(x) = 0$ , if  $C$  is a 0-certificate

and  $C$  is consistent with  $x$  (i.e.  $C(i) = 0 \Leftrightarrow x_i = 0$ ). 1-certificates for  $f$  on  $x$  are defined similarly for  $x$  such that  $f(x) = 1$ .

Let  $C_x(f)$  be the size of a smallest  $f(x)$ -certificate for  $f$  on  $x$ . Then the 0-certificate complexity of  $f$  is  $C^{(0)}(f) = \max_{x, f(x)=0} C_x(f)$ , and similarly the 1-certificate complexity of  $f$  is  $C^{(1)}(f) = \max_{x, f(x)=1} C_x(f)$ . The certificate complexity of  $f$  is  $C(f) = \max\{C^{(0)}(f), C^{(1)}(f)\}$ .

The idea behind certificates is that if the algorithm knows the bits  $x_i$ ,  $i \in S$ , it can be sure that  $f(x) = 0$  or  $f(x) = 1$ , respectively. For example, consider the OR function on  $n$  bits. For any  $x$  such that  $\text{OR}(x) = 1$ , knowing that any bit  $x_i = 1$  suffices to certify that  $\text{OR}(x) = 1$ , so  $C^{(1)}(\text{OR}) = 1$ . However, if  $x = 0^n$ , in order to certify that  $f(x) = 0$  the algorithm needs to know all  $n$  bits of the input. Hence  $C^{(0)}(\text{OR}) = n$ , so  $C(\text{OR}) = n$ .

Roughly speaking, the class of functions  $f$  such that  $D(f) = \text{polylog}(n)$  is the decision tree analogue of the time complexity class  $\text{P}$ ; similarly, functions  $f$  with  $C^{(1)}(f) = \text{polylog}(n)$ ,  $C^{(0)}(f) = \text{polylog}(n)$  correspond to  $\text{NP}$  and  $\text{co-NP}$  respectively. Recall that, in the world of time complexity, it is unknown whether  $\text{P} = \text{NP} \cap \text{co-NP}$ . In the simplified world of decision tree complexity, we can actually prove a result analogous to  $\text{P} = \text{NP} \cap \text{co-NP}$ .

**Theorem 13.4.** *For any  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ ,  $D(f) \leq C^{(0)}(f)C^{(1)}(f) \leq C(f)^2$ .*

In order to prove Theorem 13.4, we will need the following lemma.

**Lemma 13.5.** *Let  $C_0 : S \rightarrow \{0, 1\}$ ,  $C_1 : T \rightarrow \{0, 1\}$  be, respectively, a 0-certificate and a 1-certificate for a boolean function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ . Then there exists  $i \in S \cap T$  such that  $C_0(i) \neq C_1(i)$ .*

*Proof.* Assume the contrary: for all  $i \in S \cap T$ ,  $C_0(i) = C_1(i)$ . Then consider the bit-string  $x$  which assigns  $x_i = C_0(i)$  for all  $i \in S$ ,  $x_i = C_1(i)$  for all  $i \in T$ , and  $x_i = 0$  for  $i \notin S \cup T$ . This  $x$  is consistent with both  $C_0$  and  $C_1$ , contradicting the definition of certificates.  $\square$

*Proof of Theorem 13.4.* Let  $c = C^{(0)}(f)C^{(1)}(f)$ . The proof is by induction on  $c$ . For  $c = 1$ , by Lemma 13.5  $f$  depends only on one input bit  $x_i$ , so we must have  $f(x) = x_i$  or  $f(x) = \neg x_i$ . In either case  $D(f) = 1$ . So assume that  $C^{(0)}(f) = r \geq 2$  and let  $C : S \rightarrow \{0, 1\}$  be a minimal 0-certificate, for some  $S$  of size  $r$ . Our decision tree for  $f$  on input  $x$  will begin by querying all the variables in  $S$ . If the answers were consistent with  $C$ , then  $f(x) = 0$ , so we can stop. Otherwise, we recursively compute the function  $g(x')$  formed by fixing the bits of  $x$  consistent with the responses to the queries. By Lemma 13.5,  $S$  intersects every  $T$  such that  $C' : T \rightarrow \{0, 1\}$  is a 1-certificate of  $g$ . This implies that we can save one query when querying any 1-certificate  $C'$  of  $g$  (as we already know one bit of  $C'$ ). Hence, as  $C^{(0)}(g) \leq C^{(0)}(f)$  and  $C^{(1)}(g) \leq C^{(1)}(f)$ , by the inductive hypothesis

$$D(f) \leq r + C^{(0)}(g)(C^{(1)}(g) - 1) \leq r + r(C^{(1)}(f) - 1) \leq C^{(0)}(f)C^{(1)}(f),$$

which is the theorem.  $\square$

## 13.2 Randomised decision tree complexity

Also by analogy with the model of time complexity, we can generalise the decision tree model by adding the ability to use randomness. A probabilistic decision tree can be defined as follows. At each step, rather than choosing a variable to query deterministically based on previous queries, the algorithm can choose a variable to query based on a string of random bits, as well as the answers to previous queries. An alternative way of arriving at the same definition is to define a



probabilistic decision tree as a probability distribution over deterministic decision trees. The depth of a probabilistic tree  $T$  is defined as the worst-case expected number of variables queried. That is, the maximum over all  $x$  of the expected number of variables queried by  $T$  on input  $x$ .

For  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , we then let  $R(f)$  denote the randomised decision tree complexity of  $f$  – i.e. the minimal depth of  $T$ , over all randomised decision trees  $T$  that compute  $f$ . As with deterministic decision trees, we insist that  $T$  computes  $f$  with certainty. This restriction may appear to imply that no saving in queries can be made over deterministic trees, as (unlike the complexity class BPP, for instance) probabilistic trees must still always be correct on every input. However, this is not the case: we now give an example for which probabilistic trees significantly outperform deterministic trees.

Let  $\text{NAND}_n : \{0, 1\}^{2^n} \rightarrow \{0, 1\}$  denote the  $n$ -level NAND tree function, which is the function defined recursively by  $\text{NAND}_1(x_1, x_2) = \neg(x_1 \wedge x_2)$ , and

$$\text{NAND}_n(x_1, \dots, x_{2^n}) = \neg(\text{NAND}_{n-1}(x_1, \dots, x_{2^{n-1}}) \wedge \text{NAND}_{n-1}(x_{2^{n-1}+1}, \dots, x_{2^n})).$$

**Theorem 13.6.**  $D(\text{NAND}_n) = 2^n$ , but  $R(\text{NAND}_n) = O\left(\left(\frac{1+\sqrt{33}}{4}\right)^n\right)$ .  
Hence  $R(\text{NAND}_n) = O(D(\text{NAND}_n)^{0.753\dots})$ .

*Proof.* The claim that  $D(\text{NAND}_n) = 2^n$  follows from  $\deg(\text{NAND}_n) = 2^n$ , which can be proven by induction. For the second claim, consider the following randomised recursive algorithm  $\mathcal{A}_n$  for computing  $\text{NAND}_n$ . Beginning at the root, choose one of the two children at random and evaluate its subtree using  $\mathcal{A}_{n-1}$ . If it evaluates to 0, then return 1. Otherwise, evaluate the other subtree using  $\mathcal{A}_{n-1}$ .

We now analyse the expected number of queries made by this algorithm. Let  $\alpha_0(n)$  (resp.  $\alpha_1(n)$ ) denote the maximum expected number of queries over all inputs that evaluate to 0 (resp. 1). It is obvious that

$$\alpha_0(n) \leq 2\alpha_1(n-1)$$

(for inputs which evaluate to 0, both the subtrees must evaluate to 1). On the other hand, for any  $x$ , if  $\text{NAND}_n(x) = 1$ , then with probability at least  $1/2$  we evaluate only one of the subtrees. So

$$\alpha_1(n) \leq \alpha_0(n-1) + \frac{1}{2}\alpha_1(n-1).$$

So  $\alpha_1(n) \leq 2\alpha_1(n-2) + \frac{1}{2}\alpha_1(n-1)$ . Solving this recurrence gives  $\alpha_1(n) \leq \left(\frac{1+\sqrt{33}}{4}\right)^n$ ;  $\alpha_0(n)$  then obeys the same bound, up to a constant factor.  $\square$

The NAND tree example may appear somewhat contrived. However, it is essentially equivalent to a natural model for the analysis of two-player games. Consider a game like chess where the players take it in turns to make a move. Some sequences of moves lead to victory for player 1, and others to victory for player 2. For simplicity, restrict to the special case where each player only has a choice of two moves (labelled 0 or 1) at each step. In a game with only one round, player 1 wins if either of these moves is a winning move. In a game with two rounds, player 1 wins if, for each move of player 2, there exists a move player 1 can make such that the sequence of two moves leads to victory for player 1. Continuing this process, we find that in a game with  $n$  rounds (where  $n$  is even), whether player 1 wins can be determined by evaluating an AND-OR tree with  $n$  levels (an AND-OR tree is a tree whose root is labelled with AND, and vertices at each level are alternately labelled with OR and AND; see Figure 22). The leaves of the tree each correspond to a sequence




Figure 22: An AND-OR tree corresponding to a game with a winning strategy for player 1.

of moves, and are labelled with bits which specify whether each strategy is a winning strategy for player 1. Finally, observe that AND-OR trees with  $2n$  levels are the same as NAND trees with  $n$  levels (up to negating the input and output bits). We therefore see that the randomised algorithm of Theorem 13.6 gives a more efficient way of evaluating game trees.

It is not known how large the separation between randomised and deterministic decision tree complexity can be. However, it is known that the gap can be no bigger than quadratic.

**Theorem 13.7.** *For any function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ ,  $D(f) = O(R(f)^2)$ .*

*Proof.* Follows from Theorem 13.4 by observing that  $R(f) \geq C(f)$ : if a randomised algorithm is required to succeed with certainty, the bits it queries must form a certificate for  $f$ .  $\square$

It is also possible to define a variant of randomised decision tree complexity which is perhaps closer in spirit to BPP, where the algorithm is allowed to fail with probability at most  $\epsilon$  in the worst case, for some  $\epsilon < 1/2$ ; then the model of randomised decision tree complexity discussed above corresponds to  $\epsilon = 0$ . It is known that, for constant  $\epsilon$ , there can be at most a cubic separation between this model of randomised decision tree complexity and deterministic decision tree complexity.

### 13.3 Decision trees and symmetry

We now study the question of putting lower bounds on decision tree complexity in the case where  $f$  satisfies some symmetry properties. Let  $\pi \in S_n$  be a permutation of the integers  $\{1, \dots, n\}$  and for  $x \in \{0, 1\}^n$  write  $\pi(x) = (x_{\pi(1)}, \dots, x_{\pi(n)})$ . We say that  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  is invariant under  $\pi$  if  $f(x) = f(\pi(x))$  for all  $x$ . The set of permutations under which  $f$  is invariant of course forms a subgroup of  $S_n$ . We say that a group of permutations  $\Gamma$  is transitive if, for any pair  $i, j \in [n]$ , there exists  $\pi \in \Gamma$  such that  $\pi(i) = j$ .

Intuitively, functions which are invariant under a transitive permutation group “look the same” in every direction, so they should require many queries to compute. This is formalised as the following conjecture, which has been open for several decades.

**Conjecture 13.8.** *Every non-constant monotone function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  such that  $f$  is invariant under a transitive permutation group is evasive.*

A boolean function  $f$  is said to be monotone if flipping the value of an input bit from 0 to 1 cannot change the output of the function from 1 to 0. The restriction to monotone functions in Conjecture 13.8 is necessary (below we give an example of a non-monotone function which is invariant under a transitive permutation group, but is not evasive).

Conjecture 13.8 (in fact, a stronger result than this conjecture) is known to be true given certain restrictions on the input size.

**Theorem 13.9.** *Let  $n$  be a prime power. If  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  is invariant under a transitive permutation group  $\Gamma$  and  $f(0^n) \neq f(1^n)$ ,  $f$  is evasive.*

Note that, if  $f$  is monotone and non-constant,  $f(0^n) \neq f(1^n)$ .

*Proof.* Let  $n = p^r$  for some prime  $p$ . Write  $\text{orbit}(x)$  for the orbit of  $x$  under  $\Gamma$ , i.e.  $\{y : y = \pi(x), \pi \in \Gamma\}$ . We first show that, for all  $x \notin \{0^n, 1^n\}$ ,  $|\text{orbit}(x)|$  is an integer multiple of  $p$ .

We have

$$\sum_{y \in \text{orbit}(x)} \text{wt}(y) = \sum_{y \in \text{orbit}(x)} \sum_{i=1}^n y_i = \sum_{i=1}^n \sum_{y \in \text{orbit}(x)} y_i = n \sum_{y \in \text{orbit}(x)} y_1,$$

where the last equality holds because  $\Gamma$  is transitive, so there must exist  $\pi \in \Gamma$  such that  $\pi(i) = 1$ . On the other hand, all bit-strings  $y \in \text{orbit}(x)$  have the same Hamming weight, so

$$\sum_{y \in \text{orbit}(x)} \text{wt}(y) = |\text{orbit}(x)| \text{wt}(x).$$

Thus  $|\text{orbit}(x)| \text{wt}(x)$  is a multiple of  $n = p^r$ . As  $\text{wt}(x)$  is not a multiple of  $n$  for  $x \notin \{0^n, 1^n\}$ ,  $|\text{orbit}(x)|$  is a multiple of  $p$ .

By Lemma 13.2,  $f$  is evasive if

$$\sum_{x, f(x)=1} (-1)^{\text{wt}(x)} \neq 0.$$

For each  $x$  such that  $f(x) = 1$ , consider the contribution of  $\text{orbit}(x)$  to this sum. All bit-strings  $y \in \text{orbit}(x)$  have  $f(y) = 1$ ,  $\text{wt}(y) = \text{wt}(x)$ , so

$$\sum_{y \in \text{orbit}(x)} (-1)^{\text{wt}(y)} = |\text{orbit}(x)| (-1)^{\text{wt}(x)}.$$

For  $x \notin \{0^n, 1^n\}$ , this quantity is a multiple of  $p$ . Since either  $f(0^n) = 1$  or  $f(1^n) = 1$ , but not both, the sum over all  $x$  is a multiple of  $p$ , plus or minus 1. In particular, the sum is nonzero, so  $f$  is evasive.  $\square$

A particularly interesting set of boolean functions to consider in the decision tree model are graph properties. A graph property is a function  $f : \{0, 1\}^{\binom{n}{2}} \rightarrow \{0, 1\}$ , where the input bits represent the edges in an undirected graph, such that  $f$  is unchanged under permutations of the vertices of the graph. Natural graph properties include connectivity (is every vertex connected to every other vertex by an edge?) and whether a graph has a Hamiltonian path. Most such properties which have been studied are known to be evasive. However, it is known that non-evasive graph properties do exist.

A *scorpion graph* on  $n$  vertices contains a vertex  $b$  (the body) with  $n - 2$  neighbours, a vertex  $s$  (the sting) with 1 neighbour, and a vertex  $t$  (the tail) connected to both  $b$  and  $s$  (and only them). The remaining vertices form a set  $S$  (the feet) which are all connected to  $b$ . Each pair of vertices within  $S$  may or may not have an edge between them. See Figure 23 for an illustration.

**Theorem 13.10.** *There is a deterministic algorithm which determines whether a graph  $G$  is a scorpion graph using  $O(n)$  queries.*




Figure 23: A scorpion graph.

*Proof.* Exercise. □

We therefore consider only *monotone* graph properties and would like to show that any monotone graph property is evasive (i.e. the special case of Conjecture 13.8 where  $f$  is a graph property). Theorem 13.9 does not appear to immediately apply to graph properties, as  $\binom{n}{2}$  is never a power of a prime for  $n > 3$ . However, we can use the theorem to prove the following result.

**Theorem 13.11.** *Assume  $n = 2^k$  for some integer  $k$ . Then every non-trivial monotone graph property  $f : \{0, 1\}^{\binom{n}{2}} \rightarrow \{0, 1\}$  satisfies  $D(f) \geq n^2/4$ .*

*Proof.* For any integer  $j$ , let  $G_j$  be the graph made up of  $n/2^j$  disjoint copies of a clique on  $2^j$  vertices. Thus  $G_0$  is the empty graph and  $G_k$  is the complete graph on  $n$  vertices. As  $f$  is monotone and non-trivial,  $f(G_0) = 0$  and  $f(G_k) = 1$ ; further, there is a unique  $i$  such that  $f(G_i) = 0$  but  $f(G_{i+1}) = 1$ .

Imagine we are given the promise that the input graph  $G$  has the following property: the induced subgraph on vertices  $1, \dots, n/2$  consists of  $n/2^{i+1}$  disjoint cliques, as does the induced subgraph on vertices  $n/2 + 1, \dots, n$ . This implies that the only input bits (i.e. edges) on which  $f$  depends are the  $n^2/4$  potential edges between the first  $n/2$  vertices and the second  $n/2$  vertices. Let  $g : \{0, 1\}^{n^2/4} \rightarrow \{0, 1\}$  denote the restriction of  $f$  to these edges. Observe that  $g$  is still monotone. Since  $n$  is a power of 2,  $n^2/4$  is a power of a prime.  $g(0^{n^2/4}) \neq g(1^{n^2/4})$  since  $g(0^{n^2/4}) = f(G_i) = 0$  and  $g(1^{n^2/4}) = 1$  (as this corresponds to  $G_{i+1}$  with some additional edges). Also,  $g$  is invariant under the transitive permutation group generated by swapping any two vertices either side of the bipartite cut. Thus  $g$  is evasive by Theorem 13.9, so  $D(f) \geq D(g) = n^2/4$ . □

Using similar but more technical ideas, the following result can be proven.

**Theorem 13.12.** *Let  $f : \{0, 1\}^{\binom{n}{2}} \rightarrow \{0, 1\}$  be a non-trivial monotone graph property. Then  $D(f) = \Omega(n^2)$ .*

However, even for monotone graph properties, Conjecture 13.8 remains open.

### 13.4 Notes and further reading

An excellent survey on decision tree complexity and many other measures of complexity for boolean functions is “Complexity measures and decision tree complexity” by Buhrman and de Wolf, which can easily be found online. The results discussed here concerning game trees and graph properties are due to Saks and Wigderson and Rivest and Vuillemin, respectively; the exposition here is largely based on lecture notes for the course “Concrete Models of Computation” by Jeff Erickson<sup>1</sup>.

<sup>1</sup><http://compgeom.cs.uiuc.edu/~jeffe/teaching/497/05-evasive.pdf>




Figure 24: A  $r$ -message communication protocol for computing  $f(x, y)$ .

## 14 Communication complexity

The theory of communication complexity studies the amount of communication required for two (or more) spatially separated parties to accomplish some task. While a simple model, it turns out to have numerous applications in theoretical computer science. As Arora and Barak note, “it strikes the elusive balance of being simple enough so that we can actually prove strong lower bounds, but general enough so that we can obtain important applications of these lower bounds”.

In the simplest variant of the model, we imagine we have two players (conventionally known as Alice and Bob). Each has an input picked from some set (Alice has  $x \in X$ , Bob has  $y \in Y$ ). They wish to compute some known function  $f(x, y)$ , where  $f : X \times Y \rightarrow Z$ . Unfortunately, Alice does not know Bob’s input  $y$ , and Bob does not know Alice’s input  $x$ , so they have to communicate in order to compute  $f(x, y)$ . Their goal is for Bob to output  $f(x, y)$ , perhaps with some probability of failure, using the smallest possible amount of communication. They have previously agreed on a protocol to compute  $f(x, y)$ , which proceeds as follows. Alice sends a message to Bob, based on her own input  $x$ . Based on Alice’s message and his own input  $y$ , Bob chooses a message to send to Alice. Alice then replies to Bob, and this protocol continues until one of the parties (we assume Bob) outputs  $f(x, y)$ . See Figure 24 for an illustration. Note that, as with the decision tree model, we assume that computation is free: Alice and Bob are both computationally unbounded. This is a reasonable model for distributed computational tasks where communication is much more expensive than local computation.

For a function  $f : X \times Y \rightarrow Z$ , we define the deterministic communication complexity of  $f$ ,  $D_{cc}(f)$ , as the minimum amount of total communication, measured in bits, required to compute  $f(x, y)$  with certainty for any possible pair of inputs  $x, y$ . We make the following initial simple observations.

- $D_{cc}(f) \leq \lceil \log_2 |X| \rceil$ : to compute any function  $f$ , Alice can just send her input to Bob.
- If we require each message between the parties to be one bit long, it only changes the communication complexity of  $f$  by a constant factor. We henceforth assume this restriction.
- A protocol  $\mathcal{P}$  which communicates at most  $k$  bits can thus be viewed as a binary tree with at most  $2^k$  leaves.
- If we change the model so both Alice and Bob need to know the answer  $f(x, y)$ , it only changes  $D_{cc}(f)$  by an additive  $O(\log |Z|)$  term.

$$\left( \begin{array}{cc|cc} 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ \hline 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{array} \right)$$

Figure 25: A completed protocol partitions a communication matrix into monochromatic rectangles.

As with the decision tree model, it is hopefully clear that some functions can be computed using much less communication than the trivial protocol where Alice sends all her input to Bob (or vice versa). For example, consider the function  $\text{XOR2} : \{0,1\}^n \times \{0,1\}^n \rightarrow \{0,1\}$  defined by  $\text{XOR2}(x, y) = \bigoplus_{i=1}^n (x_i \oplus y_i)$ . In other words, XOR2 is defined by taking the XOR of Alice's input with Bob's input, and then the XOR of the resulting  $n$ -bit string. Using linearity of  $\oplus$ , we have

$$\bigoplus_{i=1}^n (x_i \oplus y_i) = \left( \bigoplus_{i=1}^n x_i \right) \oplus \left( \bigoplus_{j=1}^n y_j \right).$$

Thus Alice and Bob can compute XOR2 using only one bit of communication: Alice just computes the bit  $a = \bigoplus_{i=1}^n x_i$  and sends it to Bob, who has previously computed  $b = \bigoplus_{j=1}^n y_j$ . Bob then outputs  $a \oplus b$ .

### 14.1 The rank lower bound

We would like to analyse communication protocols mathematically. Define the *communication matrix* of  $f : X \times Y \rightarrow Z$  as the  $|X| \times |Y|$  matrix  $M^f$  where  $M^f_{xy} = f(x, y)$ . The rows of  $M^f$  are indexed by Alice's inputs  $x \in X$ , and the columns by Bob's inputs  $y \in Y$ . We will analyse an arbitrary communication protocol  $\mathcal{P}$  for  $f$  in terms of  $M^f$ .

A (combinatorial) rectangle in  $M^f$  is a submatrix of  $M^f$  corresponding to entries in  $A \times B$ , where  $A \subseteq X$ ,  $B \subseteq Y$ . Note that the rows and columns of  $M^f$  are not required to be consecutive, i.e. there may be gaps. We say that  $A \times B$  is monochromatic if, for all  $x \in A$ ,  $y \in B$ ,  $M^f_{xy}$  is constant. We now observe that any protocol  $\mathcal{P}$  partitions  $M^f$  into combinatorial rectangles, which follows by induction. Assume that after some number of steps of the protocol, Alice and Bob know that the input is contained in some rectangle  $R$  (this is clearly true at the start because  $X \times Y$  is a rectangle). At the next step in the protocol, one player sends a bit to the other. The choice of this bit, which depends on that player's own input, partitions  $R$  into two rectangles  $R_0, R_1$ . When the other player receives the bit, both players know that their joint inputs are contained in either  $R_0$  or  $R_1$ .

At every stage of the protocol  $\mathcal{P}$ , each rectangle in the corresponding partition gives a subset of inputs for which the communication pattern between Alice and Bob has so far been identical. Thus, if  $\mathcal{P}$  is to compute  $f(x, y)$  successfully for all  $x$  and  $y$ , it must be the case that each rectangle in the final partition corresponding to  $\mathcal{P}$  is monochromatic. Hence, if  $\mathcal{P}$  computes  $f$  and communicates  $k$  bits in total, this gives a partition of  $M^f$  into at most  $2^k$  monochromatic rectangles, as shown in Figure 25. We will use this simple idea to prove the following lower bound on communication complexity in terms of matrix rank (over the reals).

**Theorem 14.1.** *For any function  $f : X \times Y \rightarrow \{0,1\}$ ,  $D_{cc}(f) \geq \lceil \log_2 \text{rank}(M^f) \rceil$ .*

*Proof.* Let  $\mathcal{P}$  be a protocol for  $f$  that communicates  $k$  bits. Let  $L_1$  be the set of leaves of  $\mathcal{P}$  (viewed as a tree) such that the output is 1. For each  $\ell \in L_1$ , let  $R_\ell$  be the rectangle given by the subset of the inputs that reach the leaf  $\ell$ , and define the matrix  $M^\ell$  by

$$M_{xy}^\ell = \begin{cases} 1 & \text{if } (x, y) \in R_\ell \\ 0 & \text{if } (x, y) \notin R_\ell. \end{cases}$$

As the rectangles  $R_\ell$  partition  $X \times Y$ , we have  $M^f = \sum_{\ell \in L_1} M^\ell$ . By subadditivity of the rank, this implies that

$$\text{rank}(M^f) \leq \sum_{\ell \in L_1} \text{rank}(M^\ell).$$

But for each  $\ell$ ,  $\text{rank}(M^\ell) = 1$ , so  $\text{rank}(M^f) \leq |L_1| \leq 2^k$ .  $\square$

It is a long-standing but unproven conjecture that Theorem 14.1 is close to tight.

**Conjecture 14.2** (Log-Rank Conjecture). *For any function  $f : X \times Y \rightarrow \{0, 1\}$ ,  $D_{cc}(f) \leq \text{poly}(\log \text{rank}(M^f))$ .*

We now give a couple of examples demonstrating the power of this bound.

**Theorem 14.3.** *For any set  $X$ , let  $EQ : X \times X \rightarrow \{0, 1\}$  be the equality function,*

$$EQ(x, y) = \begin{cases} 1 & \text{if } x = y \\ 0 & \text{otherwise.} \end{cases}$$

*Then  $D_{cc}(EQ) = \lceil \log_2 |X| \rceil$ .*

*Proof.* The communication matrix  $M^{EQ}$  is simply the  $|X| \times |X|$  identity matrix, which has rank  $|X|$ . For the upper bound, Alice can just send all her input to Bob.  $\square$

**Theorem 14.4.** *Let  $IP : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$  be the inner product function over  $\mathbb{F}_2$ :  $IP(x, y) = x \cdot y = \bigoplus_{i=1}^n x_i y_i$ . Then  $D_{cc}(IP) = n$ .*

In order to prove this result, it will be convenient to introduce an important set of functions: the characters of the group  $\mathbb{Z}_2^n$ . These are the  $2^n$  functions  $\chi_S : \{0, 1\}^n \rightarrow \{\pm 1\}$ ,  $S \subseteq [n]$ , defined by

$$\chi_S(x) = (-1)^{\sum_{i \in S} x_i}.$$

For example,  $\chi_\emptyset(x) = 1$  for all  $x$ , and  $\chi_{[n]}(x) = (-1)^{\text{wt}(x)}$ . These functions satisfy the following important orthogonality relations.

**Lemma 14.5.** *For any  $S, T \subseteq [n]$ ,*

$$\sum_{x \in \{0, 1\}^n} \chi_S(x) \chi_T(x) = \begin{cases} 2^n & \text{if } S = T \\ 0 & \text{otherwise.} \end{cases}$$

*Proof.* We have

$$\sum_{x \in \{0, 1\}^n} \chi_S(x) \chi_T(x) = \sum_{x \in \{0, 1\}^n} (-1)^{\sum_{i \in S} x_i + \sum_{j \in T} x_j} = \sum_{x \in \{0, 1\}^n} (-1)^{\sum_{i \in S \Delta T} x_i},$$

recalling that  $S\Delta T$  denotes the symmetric difference between the sets  $S$  and  $T$ , i.e. the set of integers  $i$  in either  $S$  or  $T$ , but not both. The lemma follows from the claim that, for any  $S \neq \emptyset$ ,

$$|\{x : \bigoplus_{i \in S} x_i = 0\}| = |\{x : \bigoplus_{i \in S} x_i = 1\}|.$$

But this claim holds because there is a bijection between these two sets: for each  $x$  such that  $\bigoplus_{i \in S} x_i = 0$ , flipping bit  $j$  of  $x$ , for arbitrary  $j \in S$ , gives  $x'$  such that  $\bigoplus_{i \in S} x'_i = 1$ .  $\square$

*Proof of Theorem 14.4.* Consider the matrix  $N$  defined by  $N_{xy} = 1 - 2M_{xy}^{\text{IP}}$ . As the all 1's matrix has rank 1,  $\text{rank}(M^{\text{IP}}) \geq \text{rank}(N) - 1$ . We have

$$N_{xy} = (-1)^{\bigoplus_{i=1}^n x_i y_i}.$$

Thus, if we identify bit-strings  $y \in \{0, 1\}^n$  with subsets  $Y \subseteq [n]$  in the obvious way ( $y_i = 1 \Leftrightarrow i \in Y$ ), we see that the rows of  $N$  are given by characters of  $\mathbb{Z}_2^n$ . By Lemma 14.5, the rows of  $N$  are all orthogonal, so  $N$  has rank  $2^n$ . Thus  $D_{cc}(\text{IP}) \geq \lceil \log_2(2^n - 1) \rceil$ , which equals  $n$  for  $n \geq 2$ . The special case  $n = 1$  can be verified directly.  $\square$

## 14.2 Randomised communication complexity

As with the decision tree model, we can define a randomised variant of communication complexity. The basic idea is that Alice and Bob's protocol  $\mathcal{P}$  can depend on a string  $r$  of uniformly random bits: at each step of the protocol, each player decides what to reply based on  $r$ , as well as their inputs and the communication so far. However, there is an additional complication as compared with decision trees: do Alice and Bob each see the other's random coin flips (i.e. the randomness  $r$  is public), or do they each not know what the other's random coin flips are (i.e.  $r$  is split into two strings  $r_A, r_B$ )? These two models could potentially give rise to different measures of complexity. The second model is called "private" to contrast with "public", but note that there is no implication that Alice's random string should be kept secret from Bob, or vice versa.

For any function  $f : X \times Y \rightarrow Z$ , we define  $R_{cc}(f)$  as the minimum amount of total communication required to compute  $f(x, y)$  with success probability at least  $2/3$  on any pair of inputs  $(x, y)$ , where Alice and Bob are each given a *private* random bit-string; we define  $R_{cc}^{\text{pub}}(f)$  similarly, but give Alice and Bob a public random bit-string. As with amplification of BPP algorithms, the model is not significantly changed in either setting if the success criterion is changed from  $2/3$  to an arbitrary constant strictly greater than  $1/2$ . Observe that the length of the random string does not feature in the calculation of complexity; that is, it can be arbitrarily long. The following example illustrates the power of (public) randomness.

**Theorem 14.6.** *Let  $EQ : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$  be the equality function on two  $n$ -bit strings. Then  $R_{cc}^{\text{pub}}(EQ) \leq 2$ .*

*Proof.* Consider the following protocol  $\mathcal{P}$ . Alice and Bob jointly choose a uniformly random string  $r \in \{0, 1\}^n$ . Alice sends Bob the single bit  $x \cdot r = \bigoplus_{i=1}^n x_i r_i$ . Bob computes  $y \cdot r$  and outputs 1 if  $x \cdot r = y \cdot r$ , and 0 otherwise. If  $x = y$ , then Bob will always output 1. On the other hand, if  $x \neq y$ , by Lemma 14.5  $\Pr_r[x \cdot r \neq y \cdot r] = 1/2$ , so Bob will output 0 with probability  $1/2$ . If the players carry out  $\mathcal{P}$  twice (using independently chosen  $r$ ), and Bob outputs 1 if and only if both runs output 1, the probability of failure is at most  $1/4$ , which is strictly less than  $1/3$ .  $\square$



We thus see that  $R_{cc}^{pub}(f)$  can be almost arbitrarily small compared with  $D_{cc}(f)$ . It is clear that for any  $f$ ,  $R_{cc}^{pub}(f) \leq R_{cc}(f)$ : given a public random string  $r$ , the players can simply choose to divide it into substrings  $r_A, r_B$  such that Alice only looks at  $r_A$ , and Bob only looks at  $r_B$ . On the other hand, we now show that  $R_{cc}^{pub}(f)$  cannot be too much smaller than  $R_{cc}(f)$ .

**Theorem 14.7** (Newman's Theorem). *For any function  $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$ ,  $R_{cc}(f) = O(R_{cc}^{pub}(f) + \log n)$ .*

*Proof.* It suffices to show that any public coin protocol  $\mathcal{P}$  can be transformed into another public coin protocol  $\mathcal{P}'$  which has the same communication complexity, and such that  $\mathcal{P}'$  uses only  $O(\log n)$  random bits. Then, to make  $\mathcal{P}'$  into a private coin protocol, Alice can just generate these random bits herself and send them to Bob before the protocol starts.

Let  $Z(x, y, r)$  be the random variable which is equal to 1 if  $\mathcal{P}$  is wrong on input  $(x, y)$  and random string  $r$  (i.e.  $\mathcal{P}$  outputs 0 when it should output 1, or vice versa), and is equal to 0 otherwise. As  $\mathcal{P}$  computes  $f$  with failure probability at most  $1/3$ ,  $\mathbb{E}_r[Z(x, y, r)] \leq 1/3$  for all  $(x, y)$ . We now, iteratively, define a protocol which uses fewer random bits.

For some integer  $t > 0$ , let  $r_1, \dots, r_t$  be arbitrary bit-strings, and define the protocol  $\mathcal{P}_{r_1, \dots, r_t}$  as follows: Alice and Bob choose  $1 \leq i \leq t$  uniformly at random, and then carry out protocol  $\mathcal{P}$  with  $r_i$  as their public random string. We claim that, for any  $\delta > 0$ , if we take  $t = O(n/\delta^2)$  there exist strings  $r_1, \dots, r_t$  such that  $\mathbb{E}_i[Z(x, y, r_i)] \leq 1/3 + \delta$ , for all  $(x, y)$ . This claim will imply the theorem: Alice and Bob use the protocol  $\mathcal{P}_{r_1, \dots, r_t}$  to compute  $f$ , which uses  $O(\log n + \log(1/\delta))$  public random bits and achieves success probability at least  $2/3 - \delta$ . Taking  $\delta$  to be a small constant, we have the theorem.

It remains to prove that, for  $t = O(n/\delta^2)$ , there exist strings  $r_1, \dots, r_t$  such that  $\mathbb{E}_i[Z(x, y, r_i)] \leq 1/3 + \delta$ , for all  $(x, y)$ . To show such a set of strings exist, we use the probabilistic method and pick the strings at random. For a particular input pair  $(x, y)$ , the probability (over strings  $r_1, \dots, r_t$ ) that the failure probability is greater than  $1/3 + \delta$  is

$$\Pr_{r_1, \dots, r_t} [\mathbb{E}_i[Z(x, y, r_i)] > 1/3 + \delta] = \Pr_{r_1, \dots, r_t} \left[ \left( \frac{1}{t} \sum_{i=1}^t Z(x, y, r_i) \right) - 1/3 > \delta \right].$$

For each  $i$ ,  $\mathbb{E}_{r_i} Z(x, y, r_i) \leq 1/3$ . Thus, applying the Chernoff bound (Corollary 9.4), we get

$$\Pr_{r_1, \dots, r_t} [\mathbb{E}_i[Z(x, y, r_i)] > 1/3 + \delta] \leq e^{-\delta^2 t/3}.$$

Choosing  $t = O(n/\delta^2)$ , this is smaller than  $2^{-2n}$ . Thus, taking a union bound over all pairs  $(x, y)$ , the probability over  $r_1, \dots, r_t$  that there exists some  $(x, y)$  that  $\mathbb{E}_i[Z(x, y, r_i)] > 1/3 + \delta$  is strictly less than 1, so there must exist *some* choices  $r_1, \dots, r_t$  such that  $\mathbb{E}_i[Z(x, y, r_i)] \leq 1/3 + \delta$ . This proves the claim.  $\square$

Combining Theorems 14.6 and 14.7 implies that there exists a private coin protocol for the equality function  $\text{EQ} : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$  which uses only  $O(\log n)$  bits of communication, but does not give an explicit description of such a protocol. We now describe an explicit private-coin protocol which still uses only  $O(\log n)$  bits of communication.

Let  $p$  be a prime number such that  $n^2 < p < 2n^2$  (it follows from number-theoretic arguments which we omit that one can always find such a  $p$ ). If Alice's input is  $a \in \{0, 1\}^n$  and Bob's input is  $b \in \{0, 1\}^n$ , we think of these as polynomials  $A, B \in \mathbb{F}_p[x]$ :

$$A(x) = a_1 + a_2x + \dots + a_nx^{n-1}, \quad B(x) = b_1 + b_2x + \dots + b_nx^{n-1}.$$

The protocol proceeds as follows: Alice picks  $r \in \mathbb{F}_p$  uniformly at random and sends Bob  $r$  and  $A(r)$ ; Bob then outputs 1 if  $A(r) = B(r)$ , and 0 otherwise. The number of bits sent is thus  $2\lceil \log_2 p \rceil = O(\log n)$ . If  $a = b$ , then  $A(x) = B(x)$  for all  $x \in \mathbb{F}_p$ , so Bob always outputs 1. If  $a \neq b$  then  $A$  and  $B$  are distinct polynomials of degree  $n - 1$ . This means that they can be equal on at most  $n - 1$  elements of  $\mathbb{F}_p$  (to see this, note that  $A - B$  is a degree  $n - 1$  polynomial, which can have at most  $n - 1$  roots). Thus the probability that Bob incorrectly outputs 1 is at most  $(n - 1)/p \leq 1/n$ .

### 14.3 Lower bounds on randomised communication complexity

It is also possible to prove lower bounds on randomised communication protocols. To do so, we introduce the model of distributional communication complexity, in which the communication protocol is deterministic, but the inputs are picked according to some probability distribution.

Consider  $f : X \times Y \rightarrow Z$  and let  $\mu$  be a probability distribution on  $X \times Y$ . The  $(\mu, \epsilon)$ -distributional communication complexity of  $f$ , written  $D_{cc, \epsilon}^\mu(f)$ , is the number of bits communicated by the best deterministic protocol that outputs  $f(x, y)$  on at least a  $1 - \epsilon$  fraction (with respect to  $\mu$ ) of inputs  $(x, y) \in X \times Y$ .

For certain distributions  $\mu$ , distributional communication complexity can be much lower than randomised communication complexity; this is unsurprising as the first is an average-case model, whereas the second is a worst-case model. However, it turns out that the worst case distributional complexity over all distributions  $\mu$  completely characterises public-coin randomised communication complexity.

**Theorem 14.8.**  $R_{cc}^{pub}(f) = \max_{\mu} D_{cc, 1/3}^\mu(f)$ .

*Proof.* We will show that, for any distribution  $\mu$ ,  $R_{cc}^{pub}(f) \geq D_{cc, 1/3}^\mu(f)$ . Consider any randomised protocol  $\mathcal{P}$  for  $f$ . As for every pair of inputs  $(x, y)$   $\mathcal{P}$  outputs the right answer with probability at least  $2/3$ , averaging over any distribution  $\mu$  on the inputs  $\mathcal{P}$  is still correct with probability at least  $2/3$ , where the probability is taken both over the players' randomness and the input. Thus there must exist some choice of random coin flips such that the protocol, fixing that choice of coin flips, succeeds with probability at least  $2/3$  with respect to  $\mu$ .

We will not show the other direction here that  $R_{cc}^{pub}(f) \leq \max_{\mu} D_{cc, 1/3}^\mu(f)$ , which can be proven using von Neumann's minimax theorem from the theory of zero-sum games.  $\square$

Theorem 14.8 gives a way of proving lower bounds on randomised protocols; one picks a suitable "hard" distribution  $\mu$  and shows that no deterministic protocol using a small amount of communication can succeed on a large fraction of the inputs, with respect to  $\mu$ . One technique for showing such a lower bound on deterministic protocols is the so-called discrepancy method.

**Definition 14.1.** Let  $f : X \times Y \rightarrow \{0, 1\}$  be a function, let  $R$  be a rectangle, and let  $\mu$  be a probability distribution on  $X \times Y$ . Define

$$\text{Disc}_{\mu}(R, f) = \left| \Pr_{\mu}[f(x, y) = 0 \text{ and } (x, y) \in R] - \Pr_{\mu}[f(x, y) = 1 \text{ and } (x, y) \in R] \right|.$$

The discrepancy of  $f$  according to  $\mu$  is

$$\text{Disc}_{\mu}(f) = \max_R \text{Disc}_{\mu}(R, f),$$

where the maximum is taken over all rectangles  $R$ .

Thus, if  $f$  has low discrepancy, any rectangle in the communication matrix  $M^f$  has roughly the same number of zeros as ones (weighted by  $\mu$ ). Intuitively, this should mean it is hard for communication protocols to determine whether they should output 0 or 1. We will now see that this intuition is indeed correct.

**Theorem 14.9.** *Let  $f : X \times Y \rightarrow \{0, 1\}$  be a function, and let  $\mu$  be a probability distribution on  $X \times Y$ , and fix  $0 < \epsilon < 1/2$ . Then  $D_{cc, 1/2-\epsilon}^\mu(f) \geq \log_2(2\epsilon / \text{Disc}_\mu(f))$ .*

*Proof.* Let  $\mathcal{P}$  be a deterministic protocol which communicates  $c$  bits and computes  $f$  with success probability at least  $1/2 + \epsilon$ , with respect to  $\mu$ . Then

$$\begin{aligned} 2\epsilon &\leq \Pr_\mu[\mathcal{P}(x, y) = f(x, y)] - \Pr_\mu[\mathcal{P}(x, y) \neq f(x, y)] \\ &= \sum_\ell \left( \Pr_\mu[\mathcal{P}(x, y) = f(x, y) \text{ and } (x, y) \in R_\ell] - \Pr_\mu[\mathcal{P}(x, y) \neq f(x, y) \text{ and } (x, y) \in R_\ell] \right), \end{aligned}$$

where we split the sum up in terms of the leaves  $\ell$  of the communication protocol, and  $R_\ell$  is the rectangle corresponding to leaf  $\ell$ . As, for each  $\ell$ ,  $\mathcal{P}(x, y)$  is constant for all  $(x, y) \in R_\ell$ , this is upper bounded by

$$\sum_\ell \left| \Pr_\mu[f(x, y) = 0 \text{ and } (x, y) \in R_\ell] - \Pr_\mu[f(x, y) = 1 \text{ and } (x, y) \in R_\ell] \right| = \sum_\ell \text{Disc}_\mu(R_\ell, f) \leq 2^c \text{Disc}_\mu(f).$$

So  $2^c \text{Disc}_\mu(f) \geq 2\epsilon$ , which is the theorem.  $\square$

As an example application of this method, we now prove a lower bound on the *randomised* communication complexity of the inner product function.

**Theorem 14.10.** *Let  $IP : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$  be the inner product function over  $\mathbb{F}_2$ :  $IP(x, y) = x \cdot y = \bigoplus_{i=1}^n x_i y_i$ . Then  $R_{cc}^{\text{pub}}(IP) \geq n/2 - O(1)$ .*

*Proof.* We will use Theorem 14.9, taking  $\mu$  to be the uniform distribution. Once again, we let  $N$  be the matrix defined by  $N_{xy} = 1 - 2M_{xy}^{\text{IP}}$ . Observe that, for any rectangle  $R = A \times B$ ,

$$\text{Disc}_\mu(R, \text{IP}) = \frac{\sum_{x \in A} \sum_{y \in B} N_{xy}}{2^{2n}} = \frac{a^T N b}{2^{2n}},$$

where we write  $a$  for the characteristic vector of  $A$ , i.e.  $a_x = 1 \Leftrightarrow x \in A$ , and  $b$  for the characteristic vector of  $B$ . We showed in Lemma 14.5 that  $N$  is an orthogonal matrix (up to a normalisation factor of  $2^{n/2}$ ). Thus

$$a^T N b \leq |a^T| |N b| = 2^{n/2} |a| |b| = 2^{n/2} |A| |B| \leq 2^{3n/2},$$

where  $|v| = (\sum_i v_i^2)^{1/2}$  is the Euclidean norm of a vector  $v$ . Thus  $\text{Disc}_\mu(\text{IP}) \leq 2^{-n/2}$ . The claim then follows from Theorem 14.9.  $\square$

## 14.4 Application: Time-space tradeoffs

We now give an application of communication complexity lower bounds to the apparently unconnected topic of time-space tradeoffs for multiple-tape Turing machines. Proving lower bounds on the time complexity of explicit functions is in general a challenging task. In the case of time-space

tradeoffs, we lower our sights to attempting to show that if the space used to compute some function is low, then the time used must be high.

The rough idea behind the connection to communication complexity is as follows. We split the input tape into two parts, corresponding to two distributed inputs. As the input tape is read-only, information can only be stored on the other tapes. Thus an algorithm which uses a small amount of time and space corresponds to an efficient protocol for computing a function of these distributed inputs. This idea can be formalised as the following theorem.

**Theorem 14.11.** *Let  $f : \{0, 1\}^n \times \{0, 1\}^n \rightarrow \{0, 1\}$  be a function. Let  $M$  be a multiple-tape Turing machine which runs in time  $T(n)$  and space  $S(n)$  on inputs of size  $3n$ , accepts all inputs in the set*

$$\{x 0^n y : x, y \in \{0, 1\}^n, f(x, y) = 1\}$$

and rejects all inputs in the set

$$\{x 0^n y : x, y \in \{0, 1\}^n, f(x, y) = 0\}.$$

Then  $D_{cc}(f) = O(T(n)S(n)/n)$ .

*Proof.* Given input  $(x, y)$ , in order to compute  $f(x, y)$  Alice and Bob simulate the operation of the Turing machine  $M$  on input  $x 0^n y$ , and output 1 or 0 depending on whether  $M$  accepts or rejects this input. Divide the input tape into “ $x$ ”, “ $0$ ” and “ $y$ ” regions of  $n$  bits each. At any point in time,  $M$ ’s head on the input tape is either in the  $x$  region (in which case Alice simulates the head’s movement), in the  $y$  region (in which case Bob simulates its movement), or in the  $0$  region. In this case, the last player whose region the head was in continues to simulate its movement until it reaches the other player’s region.

Whichever player is simulating the machine at any time knows the contents of the input tape under  $M$ ’s head. In order for the simulation to work, this player must also know everything else about  $M$ ’s current configuration. This consists of  $M$ ’s current state (a constant number of bits) and the contents of each of the used cells on the other tapes, together with  $M$ ’s head positions (in total,  $O(S(n))$  bits). Whenever the simulation switches from Alice to Bob (or vice versa), this information must be communicated. As the head on the read-only tape moves at most 1 position per time step, there are at least  $n$  steps between each pair of times when the simulation switches, so there are at most  $T(n)/n$  switches in total. The total amount of communication required is therefore  $O(T(n)S(n)/n)$ .  $\square$

Theorem 14.11 may seem somewhat esoteric, but it can be applied to some fairly natural problems. For example, consider the language  $\mathcal{PAL} \subset \{0, 1\}^*$  of *palindromes*, defined as follows:

$$\mathcal{PAL} = \{ww^R : w \in \{0, 1\}^*\},$$

where  $w^R$  is the string  $w$  in reverse order. Let  $M$  be any Turing machine that decides  $\mathcal{PAL}$ . Defining  $f(x, y) = 1$  if  $x = y^R$ , and 0 otherwise,  $M$  accepts all inputs in

$$\{x 0^n y : x, y \in \{0, 1\}^n, f(x, y) = 1\}$$

and rejects all inputs in the set

$$\{x 0^n y : x, y \in \{0, 1\}^n, f(x, y) = 0\}.$$

Observe that, for inputs of size  $n$ , by the log-rank lower bound  $D_{cc}(f) = n$ . Thus, by Theorem 14.11, if  $M$  operates in time  $T(n)$  and space  $S(n)$  on inputs of size  $m = 3n$ ,  $T(n)S(n) = \Omega(n^2)$ . That is, any multiple-tape Turing machine recognising palindromes in time  $T(n)$  and  $S(n)$  satisfies  $T(n)S(n) = \Omega(n^2)$ .

The language  $\mathcal{PAL}$  can be decided in linear time and linear space (by copying the first half of the input to a work tape and checking if the reversal of the second half matches the first half), or alternatively in quadratic time and logarithmic space (by checking for whether the  $i$ 'th character of the input matches the  $(m - i + 1)$ 'th character, for each  $i$  in turn). Thus the above bound is essentially tight. This bound may seem rather weak – but in fact no unconditional super-linear time lower bounds are known for computing *any* explicitly given “natural” family of functions in the multiple-tape Turing machine model! (Note that the languages occurring in the Time Hierarchy Theorem are not considered to be “natural”.)

## 14.5 Further reading

The model of communication complexity was first introduced by Yao in 1979, and has developed into a rich and extensive field of study. Arora-Barak (chapter 13) has a brief discussion of communication complexity, but the bible of the subject is the book “Communication Complexity” by Kushilevitz and Nisan, from which the above discussion of time-space tradeoffs in particular is taken, and which contains many other applications of these ideas.

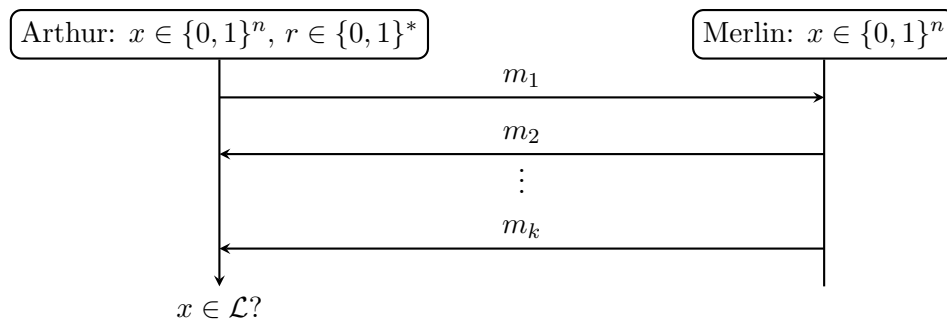


Figure 26: A  $k$ -message interactive proof for deciding  $\mathcal{L}$ .

## 15 Interactive proofs

We now link up the two concepts of communication protocols and computational complexity by studying interpretations of complexity classes in terms of interactive protocols. An alternative way to view the complexity class NP is in terms of an interaction between two players in a game: a computationally unbounded *prover* (conventionally called Merlin) and a *verifier* (called Arthur) with access to polynomial-time computation.

We imagine that Arthur is given an input  $x \in \{0, 1\}^n$ , and would like to determine whether  $x \in \mathcal{L}$ , for some language  $\mathcal{L}$  not necessarily in P, but is only able to perform deterministic polynomial-time computation. However, he also has access to a wizard (Merlin), who is computationally unbounded and can hence easily determine whether  $x \in \mathcal{L}$ . The interaction between Merlin and Arthur consists only of a single message from Merlin to Arthur. Furthermore, and unlike the previously discussed setting of communication complexity, Merlin cannot be trusted. Thus, in order for Arthur to be convinced that  $x \in \mathcal{L}$ , Merlin must send him a proof of this, which Arthur can check for himself. It is easy to see that NP is precisely the class of languages which Arthur can decide using a protocol of this form.

We now generalise this idea by allowing protocols with multiple rounds, and for Arthur to have some probability of failure. An interactive proof system with  $k$  messages is defined as follows. The input string  $x \in \{0, 1\}^n$  is known to both the prover and the verifier before the protocol starts. The verifier also has access to a string  $r$  of random bits which the prover cannot see. During the protocol, the verifier  $V$  and prover  $P$  alternately generate messages  $m_1, \dots, m_k$  to their counterpart based on any of the previous messages in the protocol, and (in the case of the verifier) the random string  $r$ . Thus the messages produced are of the form

$$m_1 = V(x, r), m_2 = P(x, m_1), m_3 = V(x, r, m_1, m_2), m_4 = P(x, m_1, m_2, m_3), \dots$$

We assume that the last message is always from the prover to the verifier, following which the verifier either accepts or rejects  $x$ . If  $k$  is even, a  $k$ -message protocol is also called a  $k/2$ -round protocol, where a round is a pair of messages (one from the verifier to the prover, and one in the opposite direction). We further assume that  $k$ ,  $r$  and the size of each message are all polynomial in  $n$ . Figure 26 illustrates such a protocol.

We are now ready to define the complexity class IP.

**Definition 15.1.** For an integer  $k \geq 1$ , we say that a language  $\mathcal{L}$  is in  $\text{IP}[k]$  if there is an interactive proof system as above, with  $k$  messages between a probabilistic polynomial-time Turing machine  $V$  and a computationally unbounded prover  $P$ , such that



Figure 27: Two isomorphic graphs.

- **(Completeness)** If  $x \in \mathcal{L}$  then there exists a prover  $P$  such that the probability that  $V$  accepts  $x$  is at least  $2/3$ ,
- **(Soundness)** If  $x \notin \mathcal{L}$  then, for all provers  $P$ , the probability that  $V$  accepts  $x$  is at most  $1/3$ ,

where probabilities are taken over the choice of  $r$ . We also write  $\text{IP} = \bigcup_{k \geq 0} \text{IP}[k]$ .

Observe that it is immediate from the definition that  $\text{BPP} = \text{IP}[0]$  and  $\text{NP} \subseteq \text{IP}[1]$ . Just as with the class  $\text{BPP}$ , we first show that this definition essentially does not depend on the choice of constants  $2/3$ ,  $1/3$ , as these can be made exponentially close to 1 and 0.

**Theorem 15.1.** *The class  $\text{IP}$  is unchanged if we replace the completeness parameter  $2/3$  by  $1 - 2^{-n^s}$ , and the soundness parameter  $1/3$  by  $2^{-n^s}$ , for any constant  $s > 0$ .*

*Proof.* The verifier simply repeats the protocol  $m$  times, and accepts if at least  $m/2$  of the runs would accept. Consider the prover which just independently repeats the strategy which each of the  $m$  sub-verifiers accepts with probability at least  $2/3$ . If  $x \in \mathcal{L}$ , then the verifier will accept  $x$  with probability at least  $1 - 2^{-\Omega(m)}$  (by a Chernoff bound, Corollary 9.4).

On the other hand, if  $x \notin \mathcal{L}$ , we need to show that every prover strategy will fail with high probability to convince the verifier that  $x \in \mathcal{L}$ . Note that this includes more complicated strategies than this independent strategy previously discussed; in particular, the prover's strategy at each stage of the protocol may depend on all the previous stages. However, at each stage the probability of the verifier accepting is still upper bounded by  $1/3$ , because the definition of soundness above holds for *all* provers. By a similar Chernoff bound argument, this implies that the verifier will accept  $x$  with probability at most  $2^{-\Omega(m)}$ .  $\square$

Although  $\text{IP}$  is at least as powerful as  $\text{NP}$ , it is not clear that  $\text{IP}$  should be any more powerful than this (e.g. compare the conjecture that  $\text{BPP} = \text{P}$ ). To build intuition, we give an example of an interactive proof for a problem which is not known to be in  $\text{NP}$ :  $\text{GRAPH NON-ISOMORPHISM}$ . This problem is the complement of the  $\text{GRAPH ISOMORPHISM}$  problem mentioned in Section 6.1. Given the description of two undirected graphs  $G$  and  $H$ , we have to accept if there is no permutation  $\pi$  of the vertices of  $G$  such that  $\pi(G) = H$  (see Figure 27). The  $\text{GRAPH ISOMORPHISM}$  problem is in  $\text{NP}$  (and thus  $\text{GRAPH NON-ISOMORPHISM}$  is in  $\text{co-NP}$ ), as a certificate of isomorphism between  $G$  and  $H$  is just the permutation  $\pi$ . However, it is not obvious how to certify that no such  $\pi$  exists, and it is not known whether  $\text{GRAPH NON-ISOMORPHISM}$  is in  $\text{NP}$ .

Nevertheless, we can give an interactive proof for  $\text{GRAPH NON-ISOMORPHISM}$ . The protocol proceeds as follows.

1. Arthur picks one of  $G$  and  $H$  at random, applies a random permutation  $\pi$  to the vertices and sends the permuted graph  $P$  to Merlin.

2. Merlin attempts to identify which of the graphs  $G$  and  $H$  was used to produce  $P$  and sends this information back to Arthur.
3. Arthur accepts if Merlin's reply is correct, and rejects otherwise.

We now prove that this protocol works. If  $G$  and  $H$  are not isomorphic, then given  $P$ , Merlin can identify whether  $P = \pi(G)$  or  $P = \pi(H)$  (he has unbounded computational power and in particular can solve GRAPH ISOMORPHISM!). Thus Merlin can reply with the correct answer and Arthur accepts with certainty. On the other hand, if  $G$  and  $H$  are isomorphic, then  $\pi(G)$  is indistinguishable from  $\pi(H)$ , so Merlin cannot guess which graph was used to produce  $P$  with probability greater than  $1/2$  (which he can achieve by just picking  $G$  or  $H$  at random). Thus Arthur incorrectly accepts isomorphic graphs with probability at most  $1/2$ . This can be reduced to  $1/4$  (to fit into the definition of IP) by performing the protocol twice and accepting only if both repetitions accept.

## 15.1 Interactive proofs and polynomial space

The following remarkable result completely characterises the interactive proof model, in terms of a rather powerful and different-looking complexity class.

**Theorem 15.2.**  $\text{IP} = \text{PSPACE}$ .

Theorem 15.2 can be split into two parts.

**Theorem 15.3.**  $\text{IP} \subseteq \text{PSPACE}$ .

*Proof.* Let  $\mathcal{L} \in \text{IP}$ , and let  $V$  be the verifier for the corresponding IP protocol for deciding  $\mathcal{L}$ . Assume the protocol has  $k$  messages. We would like to show that we can determine whether  $x \in \mathcal{L}$ , for any  $x$ , using polynomial space. It suffices to determine whether there exists a prover  $P$  such that the probability that  $V$  outputs 1 using prover  $P$  on input  $x$  is at least  $2/3$ ; if there is such a prover, then  $x \in \mathcal{L}$ , otherwise  $x \notin \mathcal{L}$ .

Let  $P(x, m_1, \dots, m_i)$  be the maximal probability that  $x$  is accepted over all provers, given that the first  $i$  messages were  $m_1, \dots, m_i$ . That is, the maximum is taken over all possible choices of prover from message  $i + 1$  onwards.  $P(x)$  is precisely the maximal probability that  $V$  accepts  $x$ , over all provers, while  $P(x, m_1, \dots, m_k)$  can be evaluated given only knowledge of the verifier's behaviour. For each  $i$ , we therefore compute  $P(x, m_1, \dots, m_{i-1})$  from  $P(x, m_1, \dots, m_i)$ . If  $m_i$  is a message from the prover to the verifier, we have

$$P(x, m_1, \dots, m_{i-1}) = \max_{m_i} P(x, m_1, \dots, m_i),$$

whereas if  $m_i$  is a message from the verifier to the prover we have

$$P(x, m_1, \dots, m_{i-1}) = \mathbb{E}_{r_i} P(x, m_1, \dots, m_i),$$

where  $r_i$  is the string of random bits used by the verifier to determine message  $m_i$ .  $P(x)$  can therefore be computed recursively using polynomial space, by a similar argument to the proof that TQBF can be solved in polynomial space.  $\square$



In order to show the reverse inclusion  $\text{PSPACE} \subseteq \text{IP}$ , it suffices to give an interactive protocol for the TQBF problem; the proof of this is somewhat technical, so we first prove a simpler result. Define the language  $3\text{-SAT}_K$  as follows:

$$3\text{-SAT}_K = \{\phi : \phi \text{ is a 3-CNF formula with exactly } K \text{ satisfying assignments}\}.$$

This is a generalisation of the co-NP-complete language  $\overline{3\text{-SAT}}$ , which is just  $3\text{-SAT}_0$ , and a restriction of the #P-complete problem  $\#3\text{-SAT}$ . We will show that  $3\text{-SAT}_K$  is in IP for any  $K$ ; in particular, this implies  $\text{co-NP} \subseteq \text{IP}$ .

## 15.2 Arithmetisation

Arithmetisation, which we briefly discussed near the start of the course, is a process establishing a link between boolean formulas and polynomials over some finite field  $\mathbb{F}$ . Given a 3-CNF formula  $\phi(x_1, \dots, x_n)$  with  $m$  clauses, introduce variables  $X_1, \dots, X_n \in \mathbb{F}$ . For each clause, write a polynomial of degree at most 3 by mapping  $x_i \mapsto (1 - X_i)$ ,  $\neg x_i \mapsto X_i$ , taking the product and subtracting the result from 1. For example,

$$x_1 \vee \neg x_4 \vee x_5 \mapsto 1 - (1 - X_1)X_4(1 - X_5),$$

and let the polynomial for the  $j$ 'th clause be denoted  $p_j(X_1, \dots, X_n)$ . For each possible assignment of 0's and 1's to  $X_1, \dots, X_n$ ,  $p_j(X_1, \dots, X_n) = 1$  if the  $j$ 'th clause is satisfied by that assignment, and  $p_j(X_1, \dots, X_n) = 0$  otherwise. Multiplying these polynomials together, we get

$$P_\phi(X_1, \dots, X_n) = \prod_{j=1}^m p_j(X_1, \dots, X_n),$$

which evaluates to 1 on satisfying assignments  $X_1, \dots, X_n$  and 0 on unsatisfying assignments. Thus  $P_\phi$  is a polynomial of degree at most  $3m$ . We now use arithmetisation to prove the following claim.

**Theorem 15.4.**  $3\text{-SAT}_K \in \text{IP}$ .

Given input  $\phi$ , we begin by constructing the polynomial  $P_\phi$ . The number of satisfying assignments of  $\phi$  is precisely

$$N_\phi := \sum_{X_1 \in \{0,1\}} \sum_{X_2 \in \{0,1\}} \cdots \sum_{X_n \in \{0,1\}} P_\phi(X_1, \dots, X_n).$$

The prover wishes to convince the verifier that  $N_\phi = K$ . The prover first sends the verifier a prime  $p$  such that  $2^n < p \leq 2^{2n}$ ; the verifier can check that  $p$  is prime using a polynomial-time primality test (e.g. the Miller-Rabin test discussed previously, or a deterministic test for primality). From now on, all computations are done over the field  $\mathbb{F}_p$ . As  $0 \leq N_\phi < p$ ,  $N_\phi = 0$  if and only if  $N_\phi \equiv 0$  modulo  $p$ .

We now give a general protocol for testing the following claim: given a prime  $p$ , an element  $K \in \mathbb{F}_p$ , and a degree  $d$  polynomial  $g(X_1, \dots, X_n) : \mathbb{F}_p^n \rightarrow \mathbb{F}_p$ ,

$$\sum_{X_1 \in \{0,1\}} \sum_{X_2 \in \{0,1\}} \cdots \sum_{X_n \in \{0,1\}} g(X_1, \dots, X_n) = K. \tag{1}$$

We assume that the verifier can evaluate  $g(X_1, \dots, X_n)$  for any  $X_1, \dots, X_n$  in time  $\text{poly}(n)$  (this is certainly true for  $g = P_\phi$ ). For each assignment of values  $b_2, \dots, b_n \in \{0, 1\}$  to  $X_2, \dots, X_n$ ,  $g(X_1, b_2, \dots, b_n)$  is a polynomial of degree at most  $d$  in the single variable  $X_1$ . Thus

$$h(X_1) := \sum_{b_2 \in \{0,1\}} \cdots \sum_{b_n \in \{0,1\}} g(X_1, b_2, \dots, b_n)$$

is also a degree  $d$  polynomial, and if claim (1) holds, then  $h(0) + h(1) = K$ . In order to check this claim, we perform the following recursive protocol.

- If  $n = 1$ , the verifier checks that  $g(0) + g(1) = K$ . If so, he accepts; otherwise, he rejects. If  $n \geq 2$ , the verifier asks the prover to send a description of the polynomial  $h(X_1)$ .
- The prover sends a polynomial  $\tilde{h}(X_1)$  (if the prover is honest,  $\tilde{h}(X_1) = h(X_1)$ ).
- The verifier rejects if  $\tilde{h}(0) + \tilde{h}(1) \neq K$ . Otherwise, he picks  $r \in \mathbb{F}_p$  uniformly at random, and uses the same protocol recursively to check that

$$\sum_{X_2 \in \{0,1\}} \cdots \sum_{X_n \in \{0,1\}} g(r, X_2, \dots, X_n) = \tilde{h}(r).$$

**Lemma 15.5.** *If (1) is false, then the verifier rejects with probability at least  $(1 - d/p)^n$ .*

*Proof.* Assuming that (1) is false, we prove the lemma by induction on  $n$ . For the base case  $n = 1$ , if  $g(0) + g(1) \neq K$  the verifier rejects with certainty. In the first round, the prover is supposed to return  $h(X_1)$ . If he does so, then as  $h(0) + h(1) \neq K$ , the verifier rejects with certainty. So assume  $\tilde{h}(X_1) \neq h(X_1)$ . Then  $\tilde{h}(X_1) - h(X_1)$  is a nonzero polynomial of degree at most  $d$ , and hence has at most  $d$  roots. Thus there are at most  $d$  values  $r$  such that  $\tilde{h}(r) = h(r)$ . Hence, when the verifier picks a random  $r$ ,  $\Pr[\tilde{h}(r) \neq h(r)] \geq 1 - d/p$ . If  $\tilde{h}(r) \neq h(r)$ , then at the next recursive step the prover has an incorrect claim to prove. By the inductive hypothesis, he fails to convince the verifier with probability at least  $(1 - d/p)^{n-1}$ . Hence the probability that the verifier rejects is at least  $(1 - d/p)(1 - d/p)^{n-1} = (1 - d/p)^n$ .  $\square$

This lemma implies Theorem 15.4, because we have  $(1 - d/p)^n \geq 1 - nd/p \geq 1 - 3mn/2^n$ , which is very close to 1 and in particular larger than  $2/3$ .

### 15.3 Extending the protocol to TQBF

The idea of arithmetisation can be used to prove that in fact  $\text{TQBF} \in \text{IP}$ , and hence  $\text{IP} = \text{PSPACE}$ . Given a quantified boolean formula

$$\psi = \forall x_1 \exists x_2 \forall x_3 \dots \exists x_n \phi(x_1, \dots, x_n),$$

where  $\phi$  is a 3-CNF<sup>1</sup> boolean formula, we create a polynomial  $P_\phi(X_1, \dots, X_n)$  in the same way as before. Then  $\psi$  is true if and only if

$$\prod_{X_1 \in \{0,1\}} \prod_{X_2 \in \{0,1\}} \prod_{X_3 \in \{0,1\}} \cdots \prod_{X_n \in \{0,1\}} P_\phi(X_1, \dots, X_n) = 1,$$

<sup>1</sup>TQBF remains PSPACE-hard under this restriction, though we technically did not show this in Theorem 8.3.

where we use the notation

$$\prod_{X_n \in \{0,1\}} g(X_n) := 1 - (1 - g(0))(1 - g(1))$$

for the arithmetisation of  $\exists$ . We will check this using essentially the same idea as for  $3\text{-SAT}_K$ , but will need to introduce a new technical trick. Imagine we want to verify that

$$\prod_{X_1 \in \{0,1\}} \prod_{X_2 \in \{0,1\}} \prod_{X_3 \in \{0,1\}} \cdots \prod_{X_n \in \{0,1\}} g(X_1, \dots, X_n) = 1 \quad (2)$$

for some polynomial  $g : \mathbb{F}_p^n \rightarrow \mathbb{F}_p$ . If we write

$$h(X_1) = \prod_{b_2 \in \{0,1\}} \prod_{b_3 \in \{0,1\}} \cdots \prod_{b_n \in \{0,1\}} g(X_1, b_2, \dots, b_n),$$

then (2) holds if and only if  $h(0)h(1) = 1$ , so checking this claim suffices to solve TQBF. Unfortunately, if  $h(X_1)$  is defined in the same way as before (i.e. recursively), its degree could become exponentially large; each operator  $\prod$  and  $\prod$  can double the degree of the polynomial. In particular,  $h(X_1)$  could become a polynomial of degree  $\Omega(2^n)$ , which means that the prover cannot necessarily communicate it efficiently to the verifier.

However, there is a trick which can fix this. The claim (2) only uses values  $X_i \in \{0,1\}$ , which implies that  $X_i^k = X_i$  for all  $k$ . This means that we can convert  $g$  into a multilinear polynomial  $\tilde{g}$  (i.e. a polynomial of degree at most 1 in each variable  $X_i$ ) such that  $\tilde{g}(X_1, \dots, X_n) = g(X_1, \dots, X_n)$  on all  $X_1, \dots, X_n \in \{0,1\}$ . Such polynomials have concise descriptions and can hence be communicated efficiently from the prover to the verifier.

For any polynomial  $p$ , let  $L_i(p)$  be the polynomial defined as

$$L_i(p)(X_1, \dots, X_n) = X_i p(X_1, \dots, X_{i-1}, 1, X_{i+1}, \dots, X_n) + (1 - X_i) p(X_1, \dots, X_{i-1}, 0, X_{i+1}, \dots, X_n).$$

It is easy to see that  $L_i(p)$  is linear in  $X_i$  and  $L_i(p)(X_1, \dots, X_n) = p(X_1, \dots, X_n)$  for  $X_1, \dots, X_n \in \{0,1\}$ . Thus, in particular,

$$L_1 L_2 \dots L_n g(X_1, \dots, X_n)$$

is multilinear for any polynomial  $g$ . So claim (2) is equivalent to

$$\prod_{X_1 \in \{0,1\}} L_1 \prod_{X_2 \in \{0,1\}} L_1 L_2 \prod_{X_3 \in \{0,1\}} L_1 L_2 L_3 \cdots \prod_{X_n \in \{0,1\}} L_1 L_2 \dots L_n g(X_1, \dots, X_n) = 1. \quad (3)$$

This expression is of size  $O(n^2)$ . As before, we now prove that claim (3) can be tested by the verifier using a recursive procedure. Assume that, for some polynomial  $g(X_1, \dots, X_k)$  and any  $b_1, \dots, b_k, K$ , the prover can convince the verifier that  $g(b_1, \dots, b_k) = K$  with probability 1 if this is true, and with probability at most  $\epsilon$  if this is false. Then let  $h(X_1, \dots, X_k)$  be the polynomial obtained by taking

$$h(X_1, \dots, X_k) = \mathcal{O}g(X_1, \dots, X_k),$$

where  $\mathcal{O}$  is any of the three operations  $\prod_{X_i \in \{0,1\}}$ ,  $\prod_{X_i \in \{0,1\}}$ , or  $L_i$ , for some  $i$ . Hence  $h$  depends on at most  $k - 1$  variables in the first two cases, and at most  $k$  in the third. Let  $d$  be an upper bound on the degree of  $h$  with respect to  $X_i$ ; for us,  $d \leq 3m$ . We want to be able to check that  $h(b_1, \dots, b_k) = K$  for any  $b_1, \dots, b_k, K$ ; we will show that if this is not true, the prover can convince the verifier that it is true with probability at most  $\epsilon + d/p$ . Assuming  $i = 1$  without loss of generality, we have three cases to check.

- $\mathcal{O} = \prod_{X_1 \in \{0,1\}}$ . The same as in the previous protocol: the prover sends the verifier a polynomial  $\tilde{h}(X_1)$  which is supposed to satisfy  $\tilde{h}(X_1) = g(X_1, b_2, \dots, b_k)$ . The verifier checks whether  $\tilde{h}(0)\tilde{h}(1) = K$ . If not, he rejects; if so, he picks  $r \in \mathbb{F}_p$  at random and asks the prover to prove  $\tilde{h}(r) = g(r, b_2, \dots, b_k)$ .
- $\mathcal{O} = \prod_{X_1 \in \{0,1\}}$ . Essentially the same, but testing whether  $1 - (1 - \tilde{h}(0))(1 - \tilde{h}(1)) = K$ .
- $\mathcal{O} = L_1$ . Once again the prover sends the verifier  $\tilde{h}(X_1)$  which is supposed to satisfy  $\tilde{h}(X_1) = g(X_1, b_2, \dots, b_k)$ . This time, the verifier picks  $r \in \mathbb{F}_p$  at random and checks that  $r\tilde{h}(1) + (1 - r)\tilde{h}(0) = K$ . If not, the verifier rejects; if so, he picks  $r' \in \mathbb{F}_p$  at random and asks the prover to prove that  $\tilde{h}(r') = g(r', b_2, \dots, b_k)$ .

The proof of correctness is now essentially the same as the remainder of the proof of Theorem 15.4.

## 15.4 Modifications to the interactive proof model

We have seen that allowing interaction and a very small probability of failure gives us the ability to solve any problem in the class PSPACE. This result is perhaps additionally surprising because, if we change completeness and soundness by only an exponentially small amount, we obtain precisely the class NP. Indeed, if we let dIP be the class obtained from Definition 15.1 by replacing the completeness parameter  $2/3$  by  $1$ , and the soundness parameter  $1/3$  by  $0$ , we have the following result.

**Theorem 15.6.**  $\text{dIP} = \text{NP}$ .

*Proof.* The inclusion  $\text{NP} \subseteq \text{dIP}$  is immediate: as discussed at the start of this section, NP is precisely dIP restricted to a single message. For the other direction, we prove that if  $\mathcal{L} \in \text{dIP}$ , in fact  $\mathcal{L} \in \text{NP}$ . If  $V$  is the dIP verifier for  $\mathcal{L}$ , a certificate that  $x \in \mathcal{L}$  is a sequence of messages  $m_1, \dots, m_k$  between the prover and  $V$  that causes  $V$  to accept. To verify this certificate, simply check that  $V(x) = m_1$ ,  $V(x, m_1, m_2) = m_3$ , etc., and that  $V(x, m_1, \dots, m_k) = 1$ . If  $x \in \mathcal{L}$ , there exists such a sequence of messages. But if such a sequence exists, then  $x \in \mathcal{L}$ , by defining a prover  $P$  such that  $P(x, m_1) = m_2$ ,  $P(x, m_1, m_2, m_3) = m_4$ , etc.  $\square$

One could also consider an extension of the model of interactive proofs to multiple provers. This is conceivably more powerful than the class IP, because the verifier can perform consistency checks across different provers. That is, if one prover is trying to trick Arthur, he might be able to determine this by using another prover to double-check the first's answers. Indeed, if we define the class MIP in exactly the same way as IP, but allow the verifier to communicate with *two* provers rather than one, we have the following result, which is arguably even more remarkable than  $\text{IP} = \text{PSPACE}$  and we state without proof.

**Theorem 15.7.**  $\text{MIP} = \text{NEXP}$ .

It can also be shown that extending the model further, to more than two provers, does not lead to any additional power: MIP with  $k$  provers is equal to MIP with 2 provers, for any  $k = \text{poly}(n)$ .

## 15.5 Historical notes and further reading

The result that  $\text{IP} = \text{PSPACE}$  was announced on Boxing Day in 1989 by Shamir, swiftly following important prior work by Lund, Fortnow, Karloff, Nisan and others; Babai has written a nice

discussion of the history of this result<sup>1</sup>. The exposition here of the proof is essentially based on lecture notes by Jonathan Katz<sup>2</sup> and Arora-Barak chapter 8; there are also some good lecture notes by Oded Goldreich<sup>3</sup>. The result that  $MIP = NEXP$  is due to Babai, Fortnow and Lund and came a matter of weeks after the proof that  $IP = PSPACE$ .

---

<sup>1</sup><http://www.cs.princeton.edu/courses/archive/spr09/cos522/BabaiEmail.pdf>

<sup>2</sup><http://www.cs.umd.edu/~jkatz/complexity/f11/lecture19.pdf>

<sup>3</sup><http://eccc.hpi-web.de/resources/pdf/cc2.pdf>

## 16 Approximation algorithms and probabilistically checkable proofs

Thus far many of the problems we have considered have been decision problems, i.e. problems with a yes-no answer. Many problems we wish to solve in practice have a different flavour and can be understood as *optimisation* problems. For example:

- **KNAPSACK**: given a set of  $n$  *values*  $\{v_i\}$  and *weights*  $\{w_i\}$ , and a maximum weight  $M$ , maximise  $\sum_{i=1}^n v_i x_i$ , subject to  $\sum_{i=1}^n w_i x_i \leq M$ ,  $x_i \in \{0, 1\}$ .
- **TRAVELLING SALESMAN**: given a graph  $G$  whose edges are weighted with costs, find a tour (cycle) that visits all vertices exactly once and achieves the lowest cost.
- **MAX-3SAT**: given a CNF boolean formula (not necessarily satisfiable) where each clause contains at most 3 variables, find an assignment to the variables that satisfies the largest number of clauses.
- **MAX-E3SAT**: the same as MAX-3SAT, but where each clause contains *exactly* 3 variables.

The decision variants of all of these are NP-complete, so we do not expect to be able to solve them exactly in polynomial time. But often in practice we might be satisfied with a “reasonably” approximate answer. The concept of the quality of an approximation can be formalised as follows.

Every optimisation problem  $P$  has a set of *feasible solutions*  $F(x)$  for each input  $x$ , and each solution  $s \in F(x)$  has a *value*  $v(s)$ <sup>1</sup>. If  $P$  is a *maximisation* problem, the optimum value is then defined as

$$\text{OPT}(x) = \max_{s \in F(x)} v(s);$$

and if  $P$  is a *minimisation* problem,

$$\text{OPT}(x) = \min_{s \in F(x)} v(s).$$

Let  $\mathcal{A}$  be an algorithm which, given an instance  $x$ , returns a feasible solution  $\mathcal{A}(x) \in F(x)$ . If  $P$  is a maximisation problem, we say that  $\mathcal{A}$  is an  $\alpha$ -approximation algorithm if, for all  $x$ , we have

$$\mathcal{A}(x) \geq \alpha \text{OPT}(x),$$

and if  $P$  is a minimisation problem, we say that  $\mathcal{A}$  is an  $\alpha$ -approximation algorithm if, for all  $x$ , we have

$$\mathcal{A}(x) \leq \alpha \text{OPT}(x).$$

For a maximisation (resp. minimisation) problem  $P$ , the *approximation threshold* of  $P$  is the largest (resp. smallest)  $\alpha$  such that there exists a polynomial-time  $\alpha$ -approximation algorithm for  $P$ . Beware that there is no standardised terminology in this field; for example, some authors redefine  $\alpha \mapsto 1/\alpha$  for maximisation problems, so it is always at least 1.

Intuitively, the approximation threshold of a problem is how well we can expect to solve it in practice. It will turn out that different NP-complete problems can have very different approximation thresholds.

---

<sup>1</sup>This terminology makes sense if we want to maximise the value; for minimisation problems the term “cost” might be more appropriate.

## 16.1 The good: arbitrarily close approximation

**Theorem 16.1.** *There is a polynomial-time  $(1 - \epsilon)$ -approximation algorithm for KNAPSACK, for any  $\epsilon > 0$ .*

*Proof.* We first use dynamic programming to get a pretty good algorithm for solving KNAPSACK. Let  $V = \max\{v_1, \dots, v_n\}$  be the maximum value of an item. Define a  $(n + 1) \times (nV + 1)$  table  $W$  by letting  $W(i, v)$  be the minimum weight attainable by selecting some subset of the first  $i$  items, such that their total value is exactly  $v$ .  $W(0, v) = \infty$  for all  $v$ , and then

$$W(i + 1, v) = \min\{W(i, v), W(i, v - v_{i+1}) + w_{i+1}\}.$$

Once we have filled in the table, we pick the largest  $v$  such that  $W(n, v) \leq M$ . This algorithm solves KNAPSACK in time  $O(n^2V)$ . Note that this is *not* polynomial in the input size, as  $V$  might be very large.

We can make this algorithm run faster, at the cost of introducing inaccuracy, by attempting to reduce  $V$ . To do this, we simply truncate the last  $b$  bits of each value  $v_i$  (for some  $b$  to be determined) and replace them with zeroes, which can then be ignored. That is, we get new values  $v'_i = 2^b \lfloor v_i / 2^b \rfloor$ . We end up with an algorithm that runs in  $O(n^2V/2^b)$  time. How bad is the solution we get? Defining  $S$  and  $S'$  to be the original and new solutions (i.e. sets of items), we have (exercise!)

$$\sum_{i \in S} v_i \geq \sum_{i \in S'} v_i \geq \sum_{i \in S'} v'_i \geq \sum_{i \in S} v'_i \geq \sum_{i \in S} (v_i - 2^b) \geq \sum_{i \in S} v_i - n 2^b.$$

Therefore, as  $V$  is a lower bound on the value of the optimal solution, this algorithm runs in time  $O(n^2V/2^b)$  and outputs a solution that is at most a  $(1 - \epsilon)$  fraction of the optimum, where  $\epsilon = n2^b/V$ .

But this implies that, for any  $\epsilon$ , if we choose  $b$  such that  $2^b = V\epsilon/n$ , we obtain an algorithm running in  $O(n^3/\epsilon)$ , which is polynomial in  $n$  for any fixed  $\epsilon$ .  $\square$

## 16.2 The bad: arbitrary inapproximability

**Theorem 16.2.** *Unless  $P = NP$ , there is no polynomial-time  $(1 + \epsilon)$ -approximation algorithm for TRAVELLING SALESMAN, for any  $\epsilon > 0$ .*

*Proof.* For a contradiction, suppose there is such an algorithm, and call it  $\mathcal{A}$ . We will use it to construct a polynomial-time algorithm for the NP-complete HAMILTONIAN CYCLE problem.

Given a graph  $G$  with  $n$  vertices, the algorithm for HAMILTONIAN CYCLE constructs a TRAVELLING SALESMAN instance with  $n$  vertices, where the distance between vertices  $i$  and  $j$  is 1 if there is an edge in  $G$  between  $i$  and  $j$ , and  $(1 + \epsilon)n$  otherwise. We now apply our approximation algorithm  $\mathcal{A}$  to this problem.

If  $\mathcal{A}$  returns a tour of cost exactly  $n$ , then we know that  $G$  has a Hamiltonian cycle. On the other hand, if  $\mathcal{A}$  returns a tour with one or more edges of cost  $(1 + \epsilon)n$ , then the total length of the tour is strictly greater than  $(1 + \epsilon)n$ . As we have assumed that  $\mathcal{A}$  never returns a tour with cost greater than  $(1 + \epsilon)$  times the optimum cost, this means that there is no tour with cost  $n$  or less. Therefore,  $G$  does not have a Hamiltonian cycle.  $\square$

### 16.3 The ugly: approximability up to a constant factor

**Theorem 16.3.** *Unless  $P = NP$ , there is no polynomial-time  $(\frac{7}{8} + \epsilon)$ -approximation algorithm for MAX-E3SAT, for any  $\epsilon > 0$ .*

There is an easy randomised  $\frac{7}{8}$ -approximation algorithm for this problem: just pick an assignment to the variables at random. A random assignment will satisfy each clause with probability  $\frac{7}{8}$ . By linearity of expectation, this implies that if there are  $m$  clauses, on average  $\frac{7}{8}m$  clauses will be satisfied. (This algorithm does not work for the more general MAX-3SAT problem. There is in fact also a  $\frac{7}{8}$ -approximation algorithm for MAX-3SAT; this algorithm is based on semidefinite programming and its proof of correctness is considerably more complicated.) Theorem 16.3 says that one cannot do any better than this trivial algorithm, unless  $P = NP$ .

We will prove a weaker variant of this result, which can be stated as follows.

**Theorem 16.4.** *Unless  $P = NP$ , there is an  $\epsilon > 0$  such that there is no polynomial-time  $(1 - \epsilon)$ -approximation algorithm for MAX-3SAT.*

Note that we have lost the tight bound on the approximation ratio and have also generalised from MAX-E3SAT to MAX-3SAT. The proof of Theorem 16.4 depends on a recently developed, and apparently unrelated, concept in theoretical computer science: probabilistically checkable proofs (PCPs).

### 16.4 Probabilistically checkable proofs

Imagine we have a proof system where a prover wishes to convince a verifier that the verifier's input is in a language  $\mathcal{L}$ . Given an  $n$ -bit input  $x$  and a poly( $n$ )-bit proof, the verifier first looks at  $x$  and performs polynomial-time computation on  $x$ . The verifier then chooses, at random, a *constant* number of bits of the proof to read. The verifier reads these bits and performs some test on them, either accepting or rejecting the input depending on the result of the test.

The **PCP Theorem** states that any language  $\mathcal{L} \in NP$  has a proof system of this form where:

- For all  $x \in \mathcal{L}$ , there exists a proof such that the verifier accepts with probability 1.
- For all  $x \notin \mathcal{L}$ , for any proof, the verifier accepts with probability at most  $1/2$ .

In comparison with the usual characterisation of NP, we are looking at much fewer bits of the proof, but at the expense of having some probability of incorrectly identifying strings as being in  $\mathcal{L}$  that are actually not in  $\mathcal{L}$ . One can generalise the notion of PCPs by writing

$$\mathcal{L} \in \text{PCP}_{c,s}(r(n), q(n))$$

to mean that there exists a probabilistic polynomial-time algorithm  $V$  which has access to  $r(n)$  random bits, and may make  $q(n)$  queries to a proof oracle  $\pi$ , such that

- **(Completeness)** For all  $x \in \mathcal{L}$ , there exists a proof  $\pi$  such that the probability that  $V$  outputs 1 on  $\pi$  is at least  $c(n)$ .
- **(Soundness)** For all  $x \notin \mathcal{L}$  and for all proofs  $\pi$ , the probability that  $V$  outputs 1 on  $\pi$  is at most  $s(n)$ .

Then it is immediate from the definition of NP that  $NP = \text{PCP}_{1,0}(0, \text{poly}(n))$ , while the PCP Theorem is that  $NP \subseteq \text{PCP}_{1,1/2}(O(\log n), O(1))$ .



## 16.5 The MAX- $q$ CSP problem

The proof of the full PCP Theorem is technical, and beyond the scope of this course. However, we will prove two related results. First, that the PCP Theorem implies an inapproximability result for MAX-3SAT. Second, we will prove a weaker variant of the PCP Theorem where the proof is exponentially long (but the verifier still only looks at a constant number of bits). To prove the first part, we first introduce a more general problem: MAX- $q$ CSP. A  $q$ CSP (“ $q$ -ary constraint satisfaction problem”) instance is a collection of boolean functions  $\phi_i : \{0,1\}^n \rightarrow \{0,1\}$  called *constraints*, each of which depends on at most  $q$  of its input bits. The MAX- $q$ CSP problem is to find an input  $x$  that satisfies the largest possible number of these constraints. MAX-3SAT is the special case of MAX-3CSP where the constraints are only allowed to be OR functions (perhaps with negations of the input bits).

We will show that the PCP Theorem implies inapproximability of MAX- $q$ CSP, for some  $q = O(1)$ . Consider a language  $\mathcal{L} \in \text{NP}$ . Given an  $n$ -bit input  $x$ , using the proof whose existence is guaranteed by the PCP Theorem, we will construct an instance  $\phi$  of MAX- $q$ CSP with  $m = \text{poly}(n)$  constraints, such that:

- If  $x \in \mathcal{L}$ ,  $\phi$  is satisfiable;
- If  $x \notin \mathcal{L}$ , the largest number of constraints of  $\phi$  that can be satisfied is at most  $\frac{m}{2}$ .

Consider a claimed proof that  $x \in \mathcal{L}$ , and introduce  $\text{poly}(n)$  variables  $\{y_i\}$ , with one variable corresponding to each bit of the proof. Let  $R$  be a string of random bits generated by the verifier. Then  $R$  corresponds to a set of bits of the proof to read, and a test function  $f_R$  to perform on them. For each  $R$ , we produce a constraint corresponding to  $f_R$ . This constraint depends on the variables that are read when random string  $R$  is generated, and is satisfied if and only if  $f_R$  passes. As the verifier only looks at a constant number of bits of the proof, there are at most  $O(\log n)$  random bits used, so at most  $\text{poly}(n)$  constraints will be produced.

Now, if  $x \in \mathcal{L}$ , by the PCP Theorem all the constraints are satisfied. On the other hand, if  $x \notin \mathcal{L}$ , at least half of the constraints must not be satisfied (or the probability of the verifier accepting would be greater than  $\frac{1}{2}$ ).

But this reduction implies that, if we could distinguish between satisfiable instances of MAX- $q$ CSP and those for which at most half the constraints are satisfiable, we could solve any problem in NP, and hence  $\text{P} = \text{NP}$ . This in turn means that MAX- $q$ CSP cannot be approximated within a ratio of  $\frac{1}{2}$ , unless  $\text{P} = \text{NP}$ .

## 16.6 Back to MAX-3SAT

The final step in the proof of Theorem 16.4 is to show that the result of the previous section implies inapproximability of MAX-3SAT, which we will achieve by converting the MAX- $q$ CSP instance  $\phi$  into a MAX-3SAT instance  $\phi'$ . First, using the universality of CNF formulae (Theorem 5.1), we can rewrite each of the clauses in  $\phi$  as a  $q$ -ary CNF expression with at most  $2^q$  clauses. Second, each of the clauses in this CNF expression can be rewritten as the AND of at most  $q$  clauses of length at most 3, using the standard reduction technique from SAT to 3-SAT. We therefore end up with a CNF expression  $\phi'$  with at most  $m' = q 2^q m$  clauses, each of which is of length at most 3.

Now it is clear that, if  $\phi$  is satisfiable,  $\phi'$  is also satisfiable. On the other hand, if an  $\epsilon$  fraction of the constraints of  $\phi$  are not satisfied, then at least an  $\frac{\epsilon}{q 2^q}$  fraction of the clauses in  $\phi'$  are not satisfied (as each constraint in  $\phi$  corresponds to at most  $q 2^q$  clauses in  $\phi'$ ). In particular, we see

that the ability to find an approximate solution to MAX-3SAT that is at least a  $(1 - (q2^{q+1})^{-1})$  fraction of the optimum implies the ability to find an approximate solution to MAX- $q$ CSP that is at least a  $\frac{1}{2}$  fraction of the optimum. Taking  $\epsilon = (q2^{q+1})^{-1}$ , we have proven Theorem 16.4: there is a constant  $\epsilon > 0$  such that there is no polynomial-time  $(1 - \epsilon)$ -approximation algorithm for MAX-3SAT, unless  $P = NP$ .

## 16.7 A weaker version of the PCP theorem

We now turn to proving the promised “toy version” of the PCP theorem. The result we will prove is the following.

**Theorem 16.5.** *For any language  $\mathcal{L} \in NP$ , there exists a proof system consisting of proofs of exponential length in the input size, where membership in  $\mathcal{L}$  can be decided with constant probability by reading only a constant number of bits of the proof. Formally,  $NP \subseteq \bigcup_{c>0} PCP_{1,1/2}(n^c, O(1))$ .*

This is much weaker than the best possible result that can be proven for proof systems with exponentially long proofs; in fact, it is known that  $\bigcup_{c>0} PCP_{1,1/2}(n^c, O(1)) = NEXP$ . The rough idea of the proof of Theorem 16.5 is as follows.

1. It suffices to give a probabilistically checkable proof of the above form for a single NP-complete problem. We will use a problem called QUADRATIC EQUATIONS over  $\mathbb{F}_2$ , which is defined as follows. Given a system  $S$  of  $m$  quadratic equations in  $n$  variables  $x_1, \dots, x_n$  over  $\mathbb{F}_2$ , where the  $k$ 'th equation is of the form

$$\bigoplus_{i,j=1}^n A_{ij}^{(k)} x_i x_j = b_k,$$

is there an assignment to the variables which satisfies all the equations?

2. Imagine there does exist such a satisfying assignment  $u$ , and that the the verifier has access to an oracle  $f_u : \{0, 1\}^{n^2} \rightarrow \{0, 1\}$  such that  $f_u(A) = \bigoplus_{i,j=1}^n A_{ij} u_i u_j$ . Then he can check that  $u$  is really a satisfying assignment by evaluating  $f_u$  on random linear combinations of the  $A^{(k)}$  matrices, and verifying that the answer is as expected.
3. The proof will thus consist of evaluations of  $f_u(A)$  for all possible  $A \in \{0, 1\}^{n^2}$ . Of course, the verifier cannot trust this proof and must verify that it is of the required form.
4. In order to make it possible for the form of the proof to be checked with only a constant number of queries, it is encoded using an error-correcting code with good error-detection properties.

First, we show that QUADRATIC EQUATIONS really is a good problem to use.

**Lemma 16.6.** QUADRATIC EQUATIONS is NP-complete.

*Proof.* We show that CIRCUIT SATISFIABILITY  $\leq_P$  QUADRATIC EQUATIONS. The proof is similar to the arithmetisation of SAT. Given a circuit  $C$  using  $n$  input variables  $x_1, \dots, x_n$  and  $m$  gates, introduce variables  $x_{n+1}, \dots, x_{n+m}$ , one for each of the gates. For each AND gate  $x_i = x_j \wedge x_k$ , add an equation  $x_i = x_j x_k$ ; for each OR gate  $x_i = x_j \vee x_k$ , add an equation  $x_i = 1 - (1 - x_j)(1 - x_k)$ ; for each NOT gate  $x_i = \neg x_j$ , add an equation  $x_i = 1 - x_j$ . Finally, for the output gate  $x_{n+m}$ , set  $x_{n+m} = 1$ . Then this system of equations has a solution if and only if  $C$  is satisfiable.  $\square$

We now introduce the code that the prover will use for satisfying assignments, which we call the linear encoding<sup>1</sup>. Each  $x \in \{0, 1\}^n$  will be encoded as a  $2^n$ -bit codeword given by evaluating the function  $L_x : \{0, 1\}^n \rightarrow \{0, 1\}$  defined by

$$L_x(y) = x \cdot y = \bigoplus_{i=1}^n x_i y_i$$

on each possible string  $y \in \{0, 1\}^n$ . This is a very inefficient encoding as it leads to an exponential blow-up in the size of  $x$ . A function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  is said to be linear if  $f(y \oplus z) = f(y) \oplus f(z)$ ; this terminology justifies the name of the above encoding, as follows.

**Lemma 16.7.** *For all  $x \in \{0, 1\}^n$ ,  $L_x : \{0, 1\}^n \rightarrow \{0, 1\}$  is linear. Furthermore, every linear function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  is a member of the set  $\{L_x : x \in \{0, 1\}^n\}$ .*

*Proof.* For the first part, we have

$$L_x(y \oplus z) = \bigoplus_{i=1}^n x_i (y_i \oplus z_i) = \left( \bigoplus_{i=1}^n x_i y_i \right) \oplus \left( \bigoplus_{i=1}^n x_i z_i \right) = L_x(y) \oplus L_x(z)$$

as required. For the second part, consider the set of bit-strings  $\{e^i : i \in [n]\}$  where  $e^i$  has a 1 at the  $i$ 'th position, and 0's elsewhere. If  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  is linear, then for any  $x \in \{0, 1\}^n$ ,

$$f(x) = \bigoplus_{i, x_i=1} f(e^i) = \bigoplus_{i=1}^n f(e^i) x_i = L_y(x),$$

where  $y \in \{0, 1\}^n$  is defined by  $y_i = 1 \Leftrightarrow f(e^i) = 1$ . □

As  $L_x(y) = L_y(x)$ , one can also view each function  $L_x(y)$  as encoding the values that every possible linear function  $L_y$  takes on input  $x$ .

**Lemma 16.8.** *For all  $x, y \in \{0, 1\}^n$  such that  $x \neq y$ ,  $|\{z : L_x(z) \neq L_y(z)\}| = 2^{n-1}$ .*

*Proof.* The claim follows from observing that the  $\{L_x\}$  functions are closely related to the characters  $\chi_S(x) = (-1)^{\sum_{i \in S} x_i}$  discussed previously in Section 14. Indeed, if  $s \in \{0, 1\}^n$  corresponds to  $S \subseteq [n]$  via  $s_i = 1 \Leftrightarrow i \in S$ , we have  $\chi_S(x) = (-1)^{L_s(x)}$  for all  $x$ . As (by Lemma 14.5) characters corresponding to different subsets are orthogonal, linear functions corresponding to different bit-strings  $x$  and  $y$  satisfy  $|\{z : L_x(z) \neq L_y(z)\}| = 2^{n-1}$ . □

Imagine the prover has a claimed solution  $u$ . He applies the linear encoding to  $u$ , and also to the  $n^2$ -bit string  $v$  formed by  $v_{ij} = u_i u_j$ , to give a  $(2^n + 2^{n^2})$ -bit proof  $(L_u, L_{uv})$ . (Of course, the verifier cannot be sure that the proof is really of this form, but must check it for himself.) Just as  $L_u$  encodes all linear functions of  $u$ ,  $L_v$  encodes all quadratic functions of  $u$ ; indeed, for  $A \in \{0, 1\}^{n^2}$ ,

$$L_v(A) = \bigoplus_{i, j=1}^n A_{ij} u_i u_j.$$

This will enable the verifier to check an arbitrary quadratic function of his choice.

---

<sup>1</sup>In the literature this is often called the Hadamard or Walsh-Hadamard code.

It will be convenient to treat these two bit-strings corresponding to all possible evaluations of  $L_u, L_v$  as oracles  $f : \{0, 1\}^n \rightarrow \{0, 1\}, g : \{0, 1\}^{n^2} \rightarrow \{0, 1\}$  which the verifier is allowed to query on arbitrary inputs. The verifier first determines whether  $f$  and  $g$  really are valid linear encodings of some original bit-strings. By Lemma 16.7, this is equivalent to determining whether they are linear functions. Therefore, it is natural to use the following *linearity test* to decide this.

**Definition 16.1** (Linearity test). Given access to a function  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , pick  $x, y \in \{0, 1\}^n$  at random and compute  $f(x), f(y)$  and  $f(x \oplus y)$ . Accept if  $f(x) \oplus f(y) = f(x \oplus y)$ ; otherwise reject.

As the linearity test only makes a few queries to  $f$  and  $g$ , it is intuitively clear that it will not be able to reject all nonlinear functions, but only those which are far from linear in some sense. For functions  $f, g : \{0, 1\}^n \rightarrow \{0, 1\}$ , we say that  $g$  is  $\epsilon$ -far from  $f$  if  $|\{x : f(x) \neq g(x)\}| = \epsilon 2^n$ . For a family of functions  $\mathcal{F}$ , we say that  $g$  is  $\epsilon$ -far from  $\mathcal{F}$  if  $\min_{f \in \mathcal{F}} |\{x : f(x) \neq g(x)\}| = \epsilon 2^n$ .

**Theorem 16.9.** *The linearity test accepts linear functions  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  with certainty. If  $f$  is  $\epsilon$ -far from linear, it rejects  $f$  with probability at least  $\epsilon$ .*

We will prove Theorem 16.9 at the end. This theorem implies that, if the linearity test passes for each of  $f$  and  $g$ , the verifier can assume that they are close to linear functions  $\tilde{f}, \tilde{g}$ . However, they may still differ from  $\tilde{f}, \tilde{g}$  on a small fraction of inputs. The following lemma shows that the verifier can in fact evaluate the linear functions  $\tilde{f}$  and  $\tilde{g}$  on arbitrary inputs.

**Lemma 16.10.** *If  $f : \{0, 1\}^n \rightarrow \{0, 1\}$  is  $\epsilon$ -far from some linear function  $\tilde{f}$ ,  $\tilde{f}(x)$  can be computed for any  $x \in \{0, 1\}^n$  with success probability  $1 - 2\epsilon$  using 2 queries to  $f$ .*

*Proof.* In order to compute  $\tilde{f}(x)$ , pick  $y \in \{0, 1\}^n$  uniformly at random, and compute  $f(x \oplus y) \oplus f(y)$ . If  $f$  is linear (i.e.  $f = \tilde{f}$ ), this equals  $\tilde{f}(x)$ . But if  $f$  is  $\epsilon$ -far from linear, then

$$\Pr_y[f(x \oplus y) \oplus f(y) \neq \tilde{f}(x)] \leq \Pr_y[f(x \oplus y) \neq \tilde{f}(x \oplus y)] + \Pr_y[f(y) \neq \tilde{f}(y)] \leq 2\epsilon.$$

□

So, using the procedure of Lemma 16.10, the verifier can essentially assume that he has access to the linear functions  $\tilde{f}, \tilde{g}$ .

**Lemma 16.11.** *Given access to linear functions  $\tilde{f} : \{0, 1\}^n \rightarrow \{0, 1\}, \tilde{g} : \{0, 1\}^{n^2} \rightarrow \{0, 1\}$ , there is a test which uses two queries to  $\tilde{f}$  and one to  $\tilde{g}$  and such that if  $\tilde{f} = L_u$  and  $\tilde{g} = L_{uu^T}$  for some  $u \in \{0, 1\}^n$ , the test accepts; otherwise, the test rejects with probability at least  $1/4$ .*

*Proof.* The test picks  $r, s \in \{0, 1\}^n$  uniformly at random and checks that  $\tilde{f}(r)\tilde{f}(s) = \tilde{g}(rs^T)$ . As  $\tilde{f}$  is linear, we know that  $\tilde{f} = L_u$  for some  $u$ . If  $\tilde{g} = L_{uu^T}$ , then

$$\tilde{g}(rs^T) = \bigoplus_{i,j=1}^n (uu^T)_{ij}(rs^T)_{ij} = \left( \bigoplus_{i=1}^n u_i r_i \right) \left( \bigoplus_{j=1}^n u_j s_j \right) = L_u(r)L_u(s) = \tilde{f}(r)\tilde{f}(s),$$

so the test passes. On the other hand, if  $\tilde{f} = L_u$  but  $\tilde{g} = L_B, B \neq uu^T$ , then applying Lemma 16.8 to a row on which  $B$  and  $uu^T$  differ, we have

$$\Pr_s[Bs \neq (uu^T)s] \geq 1/2,$$

where the multiplication is done over  $\mathbb{F}_2$ . Thus

$$\Pr_{r,s}[\tilde{g}(rs^T) \neq \tilde{f}(r)\tilde{f}(s)] = \Pr_{r,s}[r^T B s \neq r^T u u^T s] \geq 1/4,$$

which proves the claim.  $\square$

So, using the test of Lemma 16.11, the verifier can now assume that  $\tilde{g} = L_{uu^T}$  for some  $u \in \{0, 1\}^n$ . This implies that evaluating  $\tilde{g}(A)$  for any  $A \in \{0, 1\}^{n^2}$  gives

$$\bigoplus_{i,j=1}^n A_{ij} u_i u_j.$$

The verifier is now finally able to apply this to check that  $u$  satisfies the system of equations  $S$ , using the following procedure. He picks  $r \in \{0, 1\}^m$  uniformly at random and calculates

$$\tilde{g}\left(\bigoplus_{k=1}^m r_k A^{(k)}\right) = \bigoplus_{k=1}^m r_k \bigoplus_{i,j=1}^n A_{ij}^{(k)} u_i u_j.$$

If the answer is  $\bigoplus_{k=1}^m r_k b_k$ , the verifier accepts; otherwise, he rejects. This test will always accept if  $u$  does indeed satisfy all the equations; if there is at least one equation  $u$  does not satisfy, it will accept with probability  $1/2$  (by Lemma 16.8). This completes the protocol.

Let us calculate the parameters of the entire protocol. First, it is clear that the verifier only examines a constant number of bits of the proof (14 bits!). If there is a  $u$  that satisfies all the equations in  $S$ , the prover can use this  $u$  throughout and the verifier will accept with certainty. If there is no such  $u$ , we check that the verifier rejects with constant probability. Assume the linearity test succeeds with probability at least  $7/8$  on each of  $f$  and  $g$  (otherwise, the verifier will reject with constant probability). This implies that each of these functions is  $1/8$ -far from linear, so each computation of  $\tilde{f}(x)$  or  $\tilde{g}(x)$  for any  $x$  succeeds with probability at least  $3/4$ . Thus, if  $\tilde{g} \neq L_{uu^T}$  for some  $u$ , the test of Lemma 16.11 rejects with probability at least  $(3/4)^3(1/4)$ . On the other hand, if  $\tilde{g} = L_{uu^T}$  for some  $u$  that does not satisfy every equation in  $S$ , the last step of the protocol will detect this with probability  $1/2$ . So the verifier will reject incorrect “proofs” with at least constant probability; this probability can be boosted to  $1/2$  by repetition.

## 16.8 The linearity test

We still need to prove correctness of the linearity test (Theorem 16.9). To do so, we will use Fourier analysis on the group  $\mathbb{Z}_2^n$ , which we now briefly introduce.

For any positive integer  $n$ , recall that a character of the group  $\mathbb{Z}_2^n$  is a function  $\chi_S : \{0, 1\}^n \rightarrow \{\pm 1\}$ ,  $S \subseteq \{1, \dots, n\}$ , defined by  $\chi_S(x) = (-1)^{\sum_{i \in S} x_i}$ . Introduce an inner product on functions  $f, g : \{0, 1\}^n \rightarrow \mathbb{R}$ , defined by

$$\langle f, g \rangle = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} f(x)g(x).$$

We already proved the following (as Lemma 14.5).

**Theorem 16.12** (Plancherel’s theorem). *The functions  $\{\chi_S\}$  form an orthonormal basis with respect to the inner product  $\langle \cdot, \cdot \rangle$ .*

This implies that any function  $f : \{0, 1\}^n \rightarrow \mathbb{R}$  can be expanded in terms of characters, as follows:

$$f(x) = \sum_{S \subseteq [n]} \hat{f}(S) \chi_S(x),$$

for some real coefficients  $\hat{f}(S)$ , called the Fourier coefficients of  $f$ . These coefficients can be determined using the relation

$$\hat{f}(S) = \langle \chi_S, f \rangle = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} (-1)^{\sum_{i \in S} x_i} f(x).$$

Orthonormality further implies *Parseval's theorem*:

$$\frac{1}{2^n} \sum_{x \in \{0, 1\}^n} f(x)^2 = \langle f, f \rangle = \sum_{S \subseteq [n]} \hat{f}(S)^2.$$

We will want to apply Fourier analysis to boolean functions. It will be convenient to consider, instead of functions  $f : \{0, 1\}^n \rightarrow \{0, 1\}$ , functions  $f : \{0, 1\}^n \rightarrow \{\pm 1\}$  by mapping  $0 \mapsto 1$ ,  $1 \mapsto -1$  (we call this the  $\pm 1$  picture). In this picture, the linearity test previously discussed can be expressed as follows.

**Definition 16.2** (Linearity test,  $\pm 1$  picture). Given access to a function  $f : \{0, 1\}^n \rightarrow \{\pm 1\}$ , pick  $x, y \in \{0, 1\}^n$  at random and compute  $f(x)$ ,  $f(y)$  and  $f(x \oplus y)$ . Accept if  $f(x)f(y)f(x \oplus y) = 1$ ; otherwise reject.

Recall that our goal was to distinguish between the two cases of  $f$  being linear and far from linear. In the  $\pm 1$  picture, this is equivalent to distinguishing between  $f$  being a character  $\chi_S$ , or far from any character.

If  $f$  is a character, it is clear that the linearity test always accepts. More generally, we can write down the probability that the linearity test accepts as follows.

**Lemma 16.13.** *The probability that the linearity test accepts a function  $f : \{0, 1\}^n \rightarrow \{\pm 1\}$  is*

$$\frac{1}{2} + \frac{1}{2} \sum_{S \subseteq [n]} \hat{f}(S)^3.$$

*Proof.* As we have  $f(x)f(y)f(x \oplus y) \in \{\pm 1\}$ , the probability that the test accepts is equal to the expectation of the indicator random variable  $\text{ACC}(x, y) = \frac{1}{2} + \frac{1}{2}f(x)f(y)f(x \oplus y)$ . Thus

$$\Pr_{x, y}[\text{test accepts}] = \mathbb{E}_{x, y}[\text{ACC}(x, y)] = \frac{1}{2} + \frac{1}{2} \mathbb{E}_{x, y}[f(x)f(y)f(x \oplus y)].$$

We now expand this last term in terms of the Fourier expansion of  $f$ , giving

$$\begin{aligned} \mathbb{E}_{x, y}[f(x)f(y)f(x \oplus y)] &= \mathbb{E}_{x, y} \left[ \left( \sum_{S \subseteq [n]} \hat{f}(S) \chi_S(x) \right) \left( \sum_{T \subseteq [n]} \hat{f}(T) \chi_T(y) \right) \left( \sum_{U \subseteq [n]} \hat{f}(U) \chi_U(x \oplus y) \right) \right] \\ &= \sum_{S, T, U \subseteq [n]} \hat{f}(S) \hat{f}(T) \hat{f}(U) \mathbb{E}_{x, y}[\chi_S(x) \chi_T(y) \chi_U(x \oplus y)]. \end{aligned}$$

Now observe that characters are “doubly” linear in the following sense:

$$\begin{aligned}\chi_S(x)\chi_T(x) &= (-1)^{\sum_{i \in S} x_i} (-1)^{\sum_{j \in T} x_j} = (-1)^{\sum_{i \in S \Delta T} x_i} = \chi_{S \Delta T}(x), \\ \chi_S(x)\chi_S(y) &= (-1)^{\sum_{i \in S} x_i + y_i} = \chi_S(x \oplus y).\end{aligned}$$

Thus

$$\begin{aligned}\mathbb{E}_{x,y}[\chi_S(x)\chi_T(y)\chi_U(x \oplus y)] &= \mathbb{E}_{x,y}[\chi_S(x)\chi_T(y)\chi_U(x)\chi_U(y)] \\ &= \mathbb{E}_x[\chi_S(x)\chi_U(x)] \mathbb{E}_y[\chi_T(y)\chi_U(y)] \\ &= \mathbb{E}_x[\chi_{S \Delta U}(x)] \mathbb{E}_y[\chi_{T \Delta U}(y)] \\ &= \delta_{SU} \delta_{TU},\end{aligned}$$

where  $\delta_{AB} = 1$  if  $A = B$ , and 0 otherwise, so

$$\mathbb{E}_{x,y}[f(x)f(y)f(x \oplus y)] = \sum_{S \subseteq [n]} \hat{f}(S)^3$$

as required. □

We can now complete the proof of Theorem 16.9, which we restate here in the  $\pm 1$  picture.

**Theorem 16.14.** *The linearity test accepts linear functions  $f : \{0, 1\}^n \rightarrow \{\pm 1\}$  with certainty. If  $f$  is  $\epsilon$ -far from linear, it rejects  $f$  with probability at least  $\epsilon$ .*

*Proof.* If  $f$  is linear, then  $f = \chi_S$  for some  $S \subseteq [n]$ , so  $\hat{f}(S) = 1$  and by Lemma 16.13 the linearity test passes with certainty. More generally, if  $f$  is  $\epsilon$ -far from linear, then

$$1 - 2\epsilon = \max_{S \subseteq [n]} \langle f, \chi_S \rangle = \max_{S \subseteq [n]} \hat{f}(S).$$

By Lemma 16.13, the probability that the linearity test passes is

$$\frac{1}{2} + \frac{1}{2} \sum_{S \subseteq [n]} \hat{f}(S)^3 \leq \frac{1}{2} + \frac{1}{2} \max_{S \subseteq [n]} \hat{f}(S) \sum_{T \subseteq [n]} \hat{f}(T)^2 = \frac{1}{2} + \frac{1}{2} \max_{S \subseteq [n]} \hat{f}(S) = 1 - \epsilon,$$

where the inequality uses the fact that  $\hat{f}(S)^2 \geq 0$ , and the first equality is Parseval’s theorem. □

## 16.9 Further reading

Arora-Barak has much more information about PCPs (Chapters 11 and 22), including the proof of the tight bound for MAX-E3SAT; there are even entire lecture courses on the subject<sup>1</sup>. The original proof of the PCP Theorem was the culmination of a series of results drawing together many different concepts in theoretical computer science. Ryan O’Donnell has written a nice history of this<sup>2</sup>. An excellent introduction to the beautiful subject of Fourier analysis of boolean functions is provided by lecture notes from a course on the topic, again by Ryan O’Donnell<sup>3</sup>.

Finally, linearity is just one example of a property of boolean functions which is known to have an efficient tester. Oded Goldreich maintains a list of surveys about the now extensive field of property testing<sup>4</sup>.

<sup>1</sup><http://www.cs.washington.edu/education/courses/cse533/05au/>

<sup>2</sup><http://www.cs.washington.edu/education/courses/533/05au/pcp-history.pdf>

<sup>3</sup><http://www.cs.cmu.edu/~odonnell/boolean-analysis/>

<sup>4</sup><http://www.wisdom.weizmann.ac.il/~oded/test.html>