


MATHEMATICS OF THE IMPOSSIBLE

THE UNCHARTED COMPLEXITY OF COMPUTATION

Compiled on April 5, 2024

Emanuele “Manu” Viola



Contents

0.1	Conventions, choices, and caveats	7
1	A teaser	12
2	The alphabet of Time	14
2.1	Tape machines (TMs)	15
2.1.1	You don't need much to have it all	16
2.1.2	Time complexity, P, and EXP	17
2.2	TMs with large alphabet	19
2.2.1	The universal TM	19
2.3	Multi-tape machines (MTMs)	20
2.4	Circuits	20
2.5	Rapid-access machines (RAMs)	24
2.6	A challenge to the computability thesis	26
2.6.1	Robustness of BPP: Error reduction and tail bounds for the sum of random variables	27
2.6.2	Does randomness buy time?	29
2.6.3	Polynomial identity testing	30
2.6.4	Simulating BPP by circuits	32
2.6.5	Questions raised by randomness	33
2.7	Inclusion extend "upwards," separations downwards	34
2.8	Problems	34
2.9	Notes	35
3	The grand challenge	36
3.1	Information bottleneck: Palindromes requires quadratic time on TMs	37
3.2	Counting: impossibility results for non-explicit functions	39
3.3	Diagonalization and time hierarchy	40
3.3.1	$\text{TM-Time}(o(n \log n)) = \text{TM-Time}(n)$	43
3.4	Circuits	44
3.4.1	The circuit won't fit in the universe: Non-asymptotic, cosmological results	44
3.5	Problems	45
3.6	Notes	45

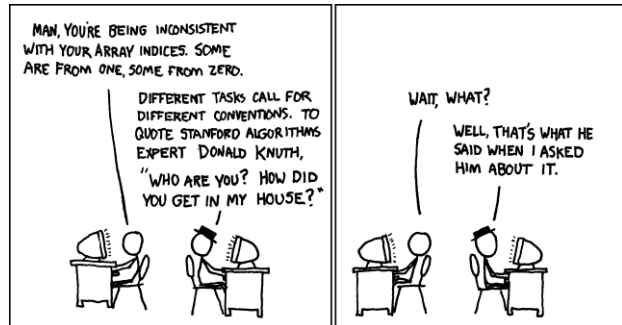
4	Reductions	46
4.1	Types of reductions	47
4.2	Reductions	48
4.2.1	Multiplication	48
4.2.2	3Sum	49
4.3	Reductions from 3Sat	50
4.3.1	3Sat to Clique	51
4.3.2	Clique to cover-by-vertexes	53
4.3.3	3Sat to Subset-Sum	53
4.3.4	3Sat to 3Color	55
4.4	Power hardness from SETH	60
4.5	Search problems	60
4.5.1	Fastest algorithm for Search-3Sat	61
4.6	Gap-SAT: The PCP theorem	62
4.7	Problems	63
4.8	Notes	63
5	Completeness: Reducing arbitrary computation	65
5.1	Nondeterministic computation	66
5.1.1	How to think of NP	67
5.2	NP-completeness	68
5.3	From RAM to 3SAT in quasi-linear time	69
5.3.1	Efficient sorting circuits: Proof of Lemma 5.1	72
5.3.2	Quasilinear-time completeness	75
5.4	Completeness in other classes	76
5.4.1	NExp completeness	76
5.4.2	Exp-completeness	77
5.5	Power from completeness	77
5.5.1	Optimization problems	77
5.5.2	NP is as easy as detecting unique solutions	78
5.6	Problems	79
5.7	Notes	80
6	Alternation	81
6.1	Does the PH collapse?	82
6.2	PH vs. AC	83
6.3	BPP in PH	85
6.4	The quantifier calculus	88
6.5	PH is a random low-degree polynomial	88
6.5.1	Back to PH	90
6.6	The power of majority	94
6.7	Problems	94
6.8	Notes	95

7	Space	96
7.1	Branching programs	98
7.2	The power of L	99
7.2.1	Arithmetic	99
7.2.2	Graphs	103
7.2.3	Linear algebra	103
7.3	Checkpoints	103
7.4	Reductions: L vs. P	105
7.4.1	Nondeterministic space	106
7.5	TiSp	109
7.6	Problems	111
7.7	Notes	111
8	Three impossibility results for 3Sat	113
8.1	Impossibility I	113
8.2	Impossibility II	114
8.3	Impossibility III	114
9	Log-depth circuits	116
9.1	The power of NC^1 : Arithmetic	117
9.2	Computing with 3 bits of memory	119
9.3	Linear-size log-depth	121
10	Constant-depth circuits	122
10.1	Threshold circuits	122
10.2	TC vs. NC^1	123
10.3	Impossibility results for AC	124
10.3.1	Impossibility results by polynomial method (a.k.a. low-degree approx- imation)	124
10.3.2	Switching lemmas	126
10.3.3	Proof of Lemma 10.2	127
10.3.4	The original switching lemma	129
10.4	ACs can sample	130
10.5	ACC	131
10.6	Notes	132
10.6.1	Myth creation: The switching lemma	132
11	Proofs	135
11.1	Static proofs	135
11.2	Proofs that yield nothing but their validity: Zero-knowledge	136
11.3	Interactive proofs	137
11.4	Interactive proofs for muggles	139
11.5	Notes	139

12 Data structures	140
12.1 Static data structures	140
12.1.1 Succinct data structures	143
12.1.2 Succincter: The trits problem	143
12.2 Dynamic data structures	147
12.3 Notes	149
13 Pseudorandomness	150
13.1 Basic PRGs	152
13.1.1 Local tests	152
13.1.2 Low-degree polynomials	153
13.1.3 Expander graphs and combinatorial rectangles: Fooling AND of sets	153
13.2 PRGs from hard functions	155
13.2.1 Turning correlation bounds into stretch: Families of sets with small intersections	158
13.2.2 Turning hardness into correlation bounds	160
13.2.3 Derandomizing the XOR lemma	162
13.2.4 Encoding the whole truth-table	163
13.2.5 Monotone amplification within NP	164
13.2.6 Proof of Theorem 13.8	164
13.3 Hardcore distributions	166
13.3.1 Proof of the hardcore-set Lemma 13.3	167
13.4 Notes	168
13.4.1 Myth creation: Polylogarithmic independence fools AC (Theorem 13.2)	169
13.5 Problems	171
14 Communication complexity	172
14.1 Two parties	172
14.1.1 The communication complexity of equality	172
14.1.2 The power of randomness	174
14.1.3 Public vs. private coins	174
14.1.4 Disjointness	174
14.1.5 Greater than	175
14.1.6 Application to TMs	175
14.1.7 Application to streaming	176
14.2 Number-on-forehead	176
14.2.1 An application to ACC	176
14.2.2 Generalized inner product is hard	177
14.2.3 Proof of Lemma 14.2	178
14.2.4 Proof of Lemma 14.3	180
14.3 Efficient protocols with logarithmically many players	181
14.3.1 The power of randomness	182
14.3.2 Pointer chasing	182

14.3.3	Sublinear communication for 3 player	184
14.4	Notes	185
15	Algebraic complexity	186
15.1	Linear transformations, rigidity, and all that	186
15.2	Computing integers	186
15.3	Univariate polynomials	188
15.4	Multivariate polynomials	188
15.4.1	VNP	189
15.5	Depth reduction in algebraic complexity	190
15.6	Completeness	191
15.7	Impossibility results for small-depth circuits	191
15.8	Algebraic TMs	191
15.9	Notes	192
16	Barriers	193
16.1	Black-box	195
16.2	Natural proofs	196
16.2.1	TMs	197
16.2.1.1	Telling subquadratic-time 1TMs from random	197
16.2.1.2	Quadratic-time 1TMs can compute pseudorandom functions	197
16.2.2	Small-depth circuits	199
16.3	Notes	199
17	I believe $P=NP$	200
18	Annotated meta-bibliography	216
19	Cryptography	218
19.1	Natural proofs	218
	Index	219

0.1 Conventions, choices, and caveats



<https://xkcd.com/163/>

I write this section before the work is complete, so some of it may change.

This book covers basic results in complexity theory and can be used for a course on the subject. At the same time, it is perhaps *sui generis* in that it contains material that, it seems, is not easy to find, possibly is even new at times, and makes somewhat unorthodox choices about topics and exposition. The book also tells a story of the quest for impossibility results, and includes some personal reflections. Some of this is discussed next.

To test your understanding of the material... this book is interspersed with mistakes, some subtle, some blatant, some not even mistakes but worrying glimpses into the author's mind. Please send all bug reports and comments to *(my five-letter last name)*@ccs.neu.edu to be included in the list of heroes.

Randomness and circuits. While randomness and circuits are everywhere in current research, and seem to be on everyone's mind, they are sometimes still relegated to later chapters, almost as an afterthought. This book starts with them right away, and attempts to weave them through the narrative.

Data structures Their study, especially negative results, squarely belongs to complexity theory. Yet data structures are strangely omitted in common textbooks. Results on data structures even tend to miss main venues for complexity theory to land instead on more algorithmic venues! We hope this book helps to revert this trend.

Algorithms & Complexity ...are of course two sides of the same coin. The rule of thumb I follow is to present algorithms that are *surprising*, *challenge our intuition of computation*, and *showcase wealth of ideas*, even though they may not be immediately deployed.

The c notation. The mathematical symbol c has a special meaning in this text. Every *occurrence* of c denotes a real number > 0 . There exist choices for these numbers such that the claims in this book are (or are meant to be) correct. This replaces, is more compact than, and is less prone to abuse than the big-Oh notation (sloppiness hides inside brackets).

Example 0.1. “For all sufficiently large n ” can be written as $n \geq c$.

“For every ϵ and all sufficiently large n ” can be written as $n \geq c_\epsilon$.

The following are correct statements:

“It is an open problem to show that some function in NP requires circuits of size cn .”

At the moment of this writing, one can replace this occurrence with 5. Note such a claim will remain true if someone proves a $6n$ lower bounds. One just needs to “recompile” the constants in this book.

“ $c > 1 + c$ ”, e.g. assign 2 to the first occurrence, 1 to the second.

“ $100n^{15} < n^c$ ”, for all large enough n . Assign $c = 16$.

The following are not true:

“ $c < 1/n$ for every n ”. No matter what we assign c to, we can pick a large enough n .

Note the assignment to c is absolute, independent of n .

More generally, when subscripted this notation indicates a function of the subscript. There exist choices for these functions such that the claims in this book are (or are meant to be) correct. Again, each occurrence can indicate a different function.

For the reader who prefers the big-Oh notation a quick and dirty fix is to replace every occurrence of c in this book with $O(1)$.

The alphabet of TMs. I define TMs with a fixed alphabet. This choice slightly simplifies the exposition (one parameter vs. two), while being more in line with common experience (it is more common experience to increase the length of a program than its alphabet). This choice affects the proof of Theorem 3.4; but the details don’t seem any worse.

Partial vs. total functions (a.k.a. on promise problems).

Recall that *promise problems offer the most direct way of formulating natural computational problems.* [...] In spite of the foregoing opinions, we adopt the convention of focusing on standard decision and search problems. [68]

I define complexity w.r.t. *partial* functions whereas most texts consider *total* functions, i.e. we consider computing functions with arbitrary domains rather than any possible string. This is sometimes called “promise problems.” This affects many things, for example the hierarchy for BPTIME (Exercise 3.4).

References and names. I decided to keep references in the main text to a minimum, just to avoid having a long list later with items “Result X is due to Y,” but relegate discussion to bibliographic notes. **Update 2024:** After consideration, I decided to remove references

in the main text altogether, to avoid equity issues with my previous plan. The new plan is being implemented.

I have also decided to not spell out names of authors, which is increasingly awkward. Central results, such as the PCP theorem, are co-authored by five or more people. But I don't mean to deprive the reader entirely of the thrill of name-splashing. So names appear in select portions which bend to the historical. **Update 2024:** Considering this: Names also appear in the index, so one can look up “Markov’s inequality” there. For format, look up “Markov’s inequality.”

For who got what award for what see [136].

Polynomial. It is customary in complexity theory to bound quantities by a polynomial, as in polynomial time, when in fact only one monomial matters. It seems to me this makes some statements cumbersome, and lends itself to confusion since polynomials with many terms are useful for many other things. I use *power* instead of polynomial, as in power time. One issue is that “power” is not an adjective. However, terminology such as “power law” is commonplace, and quite apt.

Random-access machines. “Random access” also leads to strange expressions like “randomized random-access” [14].

Reductions. Are presented as an implication. Clashing with most texts, this affects several things, for example the definition of NP-intermediate problems, see Exercise 5.3.

Exercises, problems, and questions. Exercises are interspersed within the narrative and serve as “concept check.” They are not meant to be difficult or new, though some are. Problems are collected at the end and tend to be harder and more original, though some are not. Questions are meant as *research* questions, or open problems, or challenges.

On presenting proofs. I try to present them in a “top down” fashion rather than “bottom up,” starting with the main high-level ideas and then progressively opening up details. Both styles can be abused, but it seems to me abuse of the latter is more widespread and problematic.

Summary of some terminological and not choices

Some other sources	this book	acronym
$O(1), \Omega(1)$	c	
Turing machine	tape machine	TM
random-access machine	rapid-access machine	RAM
polynomial time	power time	P
superpolynomial	superpower	
mapping reduction (sometimes)	A reduces to B in P means $B \in P \Rightarrow A \in P$	
Extended Church-Turing thesis	Power-time computability thesis	
pairwise independent	pairwise uniform	
FP, promise-P	P	
TM with any alphabet	TM with fixed alphabet	
classes have total functions	classes have partial functions	
AC^0	AC	
TC^0	TC	
P/poly	P Ckt	
$\{0, 1\}$	[2]	

The $\{0, 1\}$ notation is cumbersome for people and compilers. What I really would like is use $2 = \{0, 1\}$ as in $f : 2^n \rightarrow 2$, but I fear it's pushing it a little

Acknowledgments

Obviously, this work would not have been possible had I not had the privilege of interacting with many (most?) of the leading complexity theorists over the years. I am indebted to the students in my Spring 2023 Complexity Theory class, during which this material was first flushed out, for much useful feedback, and similarly to the students in my Spring 2024 Complexity Theory class, and its TA, Dustin Lin. I am similarly indebted to the readers of my blog, where this material was first serialized. Theorem 2.6 was sparked by Michal Koucky's talk "Journey with Eric Allender: From Turing Machines to Circuits...." The proof of Theorem 7.2 was a class project.

Thanks to Marco Genovesi for his rendering of the cover image.

Unindexed mathematical notation and abbreviations

$[i..j]$	$\{i, i + 1, i + 2, \dots, j\}$	
$[i]$	$[0..i - 1] = \{0, 1, 2, \dots, i - 1\}$	
a.k.a.	also known as	
e.g.	as an example (exempli gratia)	
i.e.	that is (id est)	
iff	if and only if	
lhs	left-hand side	
prob.	probability	
rhs	right-hand side	
r.v.	random variable	
s.t.	such that	
w.h.p.	with high prob.	
w.l.o.g.	without loss of generaliy	

Copyright by Emanuele Viola

Chapter 1

A teaser

Consider a computer with *three* bits of memory. There's also a clock, beating 1, 2, 3, ... In one clock cycle the computer can read one bit of the input and update its memory, or stop and return a value. These actions depend only on the clock, the three memory bits, and the length of the input.

Let's give a few examples of what such computer can do.

First, it can compute the And function on n bits:

```
Computing And of  $(x_1, x_2, \dots, x_n)$ 
```

```
For  $i = 1, 2, \dots$  until  $n$ 
```

```
  Read  $x_i$ 
```

```
  If  $x_i = 0$  return 0
```

```
Return 1
```

We didn't really use the memory. Let's consider a slightly more complicated example. A word is *palindrome* if it reads the same both ways, like *racecar*, *non*, *anna*, and so on. Similarly, example of palindrome bit strings are 11, 0110, and so on.

Let's show that the computer can decide if a given string is palindrome quickly, in n steps

```
Deciding if  $(x_1, x_2, \dots, x_n)$  is palindrome:
```

```
For  $i = 1, 2, \dots$  until  $i > n/2$ 
```

```
  Read  $x_i$  and write it in memory bit  $m$ 
```

```
  If  $m \neq x_{n-i}$  return 0
```

```
Return 1
```

That was easy. Now consider the Majority function on n bits, which is 1 iff the sum of the input bits is $> n/2$ and 0 otherwise. Majority, like any other function on n bits, can be computed on such a computer in time *exponential* in n . To do that, you do a pass on the input and check if it's all zero, using the program for And given above. If it is, return 0. If it is not, you do another pass now checking if it's all zero except the last bit is 1. If it is, return 0. You continue this way until you exhausted all the $2^n/2$ possible inputs with Majority equals to 0. If you never returned 0 you can now safely return 1.

As we said, this works for any function, but it's terribly inefficient. Can we do better for Majority? Can we compute it in time which is just a power of n ?

Exercise 1.1. Convince yourself that this is impossible. Hint: If you start counting bits, you'll soon run out of memory.

If you solved the exercise, you are not alone.
And yet, we will see the following shocking result:

Theorem 1.1. Majority can be computed on such a computer in time n^c .

And this is not a trick tailored to majority. Many other problems, apparently much more complicated, can also be solved in the same time.

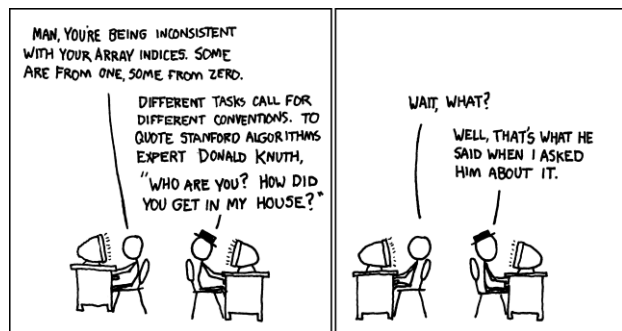
But, there's something possibly even more shocking.

Shocking situation:
It is consistent with our state of knowledge that every "textbook algorithm" can be solved in time n^c on such a computer! Nobody can disprove that. (Textbook algorithms include sorting, maxflow, dynamic programming algorithms like longest common subsequence etc., graph problems, numerical problems, etc.)

The **Shocking theorem** gives some explanation for the **Shocking situation**. It will be hard to rule out efficient programs on this model, since they are so powerful and counterintuitive. In fact, we will see later that this can be formalized. Basically, we will show that the model is so strong that it can compute functions that provably escape the reach of current mathematics... if you believe certain things, like that it's hard to factor numbers. This now enters some of the *mysticism* that surrounds complexity theory, where different beliefs and conjectures are pitted against each other in a battle for ground truth.

Chapter 2

The alphabet of Time



<https://xkcd.com/163/>

The details of the model of computation are not too important if you don't care about power differences in running times, such as the difference between solving a problem on an input of length n in time cn vs. time cn^2 . But they matter if you do.

The fundamental features of computation are two:

- **Locality.** Computation proceeds in small, local steps. Each step only depends on and affects a small amount of “data.” For example, in the grade-school algorithm for addition, each step only involves a constant number of digits.
- **Generality.** The computational process is general in that it applies to many different problems. At one extreme, we can think of a single algorithm which applies to an infinite number of inputs. This is called *uniform* computation. Or we can design algorithms that work on a finite set of inputs. This makes sense if the description of the algorithm is much smaller than the description of the inputs that can be processed by it. This setting is usually referred to as *non-uniform* computation.

Keep in mind these two principles when reading the next models.

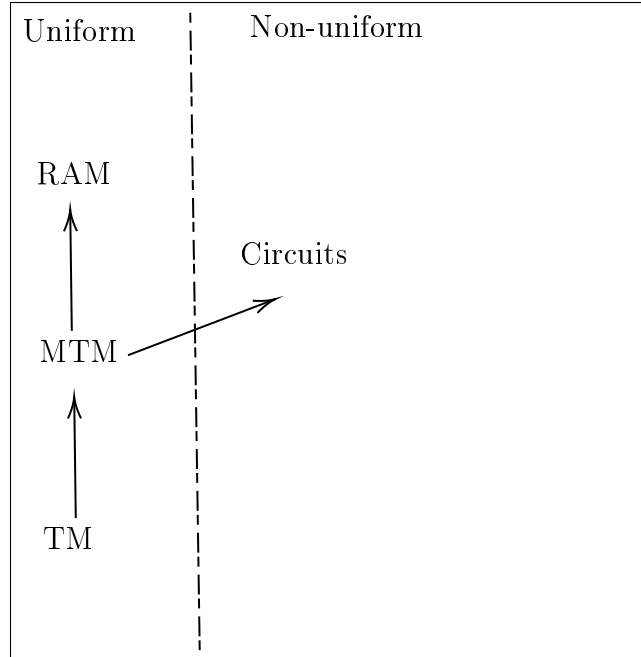


Figure 2.1: Computational models for Time. An arrow from A to B means that B can simulate A efficiently (from time t to $t \log^c t$).

2.1 Tape machines (TMs)

Tape machines are equipped with an infinite tape of cells with symbols from the *tape alphabet* A , and a *tape head* lying on exactly one cell. The machine is in one of several states, which you can think of as lines in a programming language. In one step the machine writes a symbol where the head points, changes state, and moves the head one cell to the right or left. Alternatively, it can stop. Such action depends only on the state of the machine and the tape symbol under the head.

We are interested in studying the *resources* required for computing. Several resources are of interest, like time and space. In this chapter we begin with time.

Definition 2.1. [193] A *tape machine* (TM) with s states is a map (known as *transition* or *step*)

$$\sigma : \{\underline{1}, \underline{2}, \dots, \underline{s}\} \times A \rightarrow A \times \{\text{Left}, \text{Right}, \text{Stop}\} \times \{\underline{1}, \underline{2}, \dots, \underline{s}\},$$

where $A := \{0, 1, \#, -, _ \}$ is the *tape alphabet*. The alphabet symbol $_$ is known as *blank*.

A *configuration* of a TM encodes its tape content, the position of the head on the tape, and the current state. It can be written as a triple (M, i, \underline{j}) where M maps the integers to A and specifies the tape contents, i is an integer indicating the position of the head on the tape, and \underline{j} is the state of the machine.

A configuration (μ, i, \underline{j}) *yields* $(\mu', i + 1, \underline{j}')$ if $\sigma(\underline{j}, \mu[i]) = (a, \text{Right}, \underline{j}')$ and $\mu'[i] = a$ and $\mu' = \mu$ elsewhere, and similarly it yields $(\mu', i - 1, \underline{j}')$ if $\sigma(\underline{j}, \mu[i]) = (a, \text{Left}, \underline{j}')$ and $\mu'[i] = a$ and $\mu' = \mu$ elsewhere, and finally it yields itself if $\sigma(\underline{j}, \mu[i]) = (a, \text{Stop}, \underline{j}')$.

We say that a TM computes $y \in [2]^*$ on input $x \in [2]^*$ in *time* t (or in t steps) if, starting in configuration $(\mu, 0, \underline{1})$ where $x = \mu[0]\mu[1] \cdots \mu[|x| - 1]$ and μ is blank elsewhere, it yields a sequence of t configurations where the last one is (μ, i, \underline{j}) where $\sigma(\mu[i], \underline{j})$ has a Stop instruction, and $y = \mu[i]\mu[i + 1] \cdots \mu[i + |y| - 1]$ and μ is blank elsewhere.

Describing TMs by giving the transition function quickly becomes complicated and uninformative. Instead, we give a high-level description of how the TM works. The important points to address are how the head moves, and how information is moved across the tape.

Example 2.1. On input $x \in [2]^*$ we wish to compute $x + 1$ (i.e., we think of x as an integer in binary, and increment by one). This can be accomplished by a TM with c states as follows. Move the head to the least significant bit of x . If you read a 0, write a 1, move the head to the beginning, and stop. If instead you read a 1, write a 0, move the head by one symbol, and repeat. If you reach the beginning of the input, shift the input by one symbol, append a 1, move the head to the beginning and stop.

The TM only does a constant number of passes over the input, so the running time is $c|x|$.

Example 2.2. On an input $x \in [2]^*$ we wish to decide if it has the the same number of zeros and ones. This can be done as follows. Do a pass on the input, and cross off one 0 and one 1 (by replacing them with tape symbol $\#$). If you didn't find any 0 or or 1, accept (that is, write 1 on the tape and stop). If only find a 0 but not a 1, or vice versa, reject.

Since every time we do a pass we cross at least two symbols, the running time is cn^2 .

Exercise 2.1. Describe a TM that decides if a string $x \in [2]^*$ is palindrome, and bound its running time.

Exercise 2.2. Describe a TM that on input $x \in [2]^*$ computes $n := |x|$ in binary in time $cn \log n$.

TMs can compute any function if they have sufficiently many states:

Exercise 2.3. Prove that every function $f : [2]^n \rightarrow [2]$ can be computed by a TM in time n using 2^{n+1} states.

2.1.1 You don't need much to have it all

How powerful are tape machines? Perhaps surprisingly, they are all-powerful.

Power-time computability thesis. For any “realistic” computational model C there is $d > 0$ such that: Anything that can be computed on C in time t can also be computed on a TM in time t^d .

This is a *thesis*, not a *theorem*. The meaning of “realistic” is a matter of debate, and one challenge to the thesis is discussed in section §2.6.

However, the thesis can be proved for many standard computational models, which include all modern programming languages. The proofs aren't hard. One just tediously goes through each instruction in the target model and gives a TM implementation. We prove a representative case below (Theorem 2.7) for *rapid-access machines* (RAMs), which are close to how computers operate, and from which the jump to a programming language is short.

Given the thesis, why bother with TMs? Why not just use RAMs or a programming language as our model? In fact, we will basically do that. Our default for complexity will be RAMs. However, some of the benefits of TMs remain

- TMs are easier to define – just imagine how more complicated Definition 2.1 would be were we to use a different model. Whereas for TMs we can give a short self-contained definition, for other models we have to resort to skipping details. There is also some arbitrariness in the definition of other models. What operations exactly are allowed?
- TMs allow us to pinpoint more precisely the limits of computation. Results such as Theorem ?? are easier to prove for TMs. A proof for RAM would first go by simulating RAM by a TM.
- Finally, TMs allow us to better pinpoint the limits of our knowledge about computation; we will see several examples of this.

In short, RAMs and programming languages are useful to carry computation, TMs to analyze it.

2.1.2 Time complexity, P, and EXP

We now define our first complexity classes. We are interested in solving a variety of computational tasks on TMs. So we make some remarks before the definition.

- We often need to compute *structured* objects, like tuples, graphs, matrices, etc. One can encode such objects in binary by using multiple bits. We will assume that such encodings are fixed and allow ourselves to speak of such structures objects. For example, we can encode a tuple (x_1, x_2, \dots, x_t) where $x_i \in [2]^*$ by repeating each bit in each x_i twice, and separate elements with 01.
- We can view machines as *computing functions*, or *solving problems*, or *deciding sets*, or *deciding languages*. These are all equivalent notions. For example, for a set A , the problem of deciding if an input x belongs to A , written $x \in A$, is equivalent to computing the boolean characteristic function f_A which outputs 1 if the input belongs to A , and 0 otherwise. We will use this terminology interchangeably. In general, “computing a function” is more appropriate terminology when the function is not boolean.
- We allow *partial functions*, i.e., functions with a domain X that is a strict subset of $[2]^*$, as opposed to *total functions* which are defined over $[2]^*$ or $[2]^n$. Partial functions are a natural choice for many problems, cf. discussion in section §0.1.

- We measure the running time of the machine in terms of the *input length*, usually denoted n . Input length can be a coarse measure: it is often natural to express the running time in terms of other parameters (for example, the time to factor a number could be better expressed in terms of the number of factors of the input, rather than its bit length). However for most of the discussion this coarse measure suffices, and we will discuss explicitly when it does not.
- We allow non-boolean outputs. However the running time is still only measured in terms of the input. (Another option which sometimes makes sense, it to bound the time in terms of the output length as well, which allows us to speak meaningfully of computing functions with very large outputs, such as exponentiation.)
- More generally, we are interested in computing not just functions but *relations*. That is, given an input x we wish to compute some y that belongs to a *set* $f(x)$. For example, the problem at hand might have more than one solution, and we just want to compute any of them.
- We are only interested in sufficiently large n , because one can always hard-wire solutions for inputs of fixed size, see Exercise 2.3. This allows us to speak of running times like $t(n) = n^2/1000$ without worrying that it is not suitable when n is small (for example, $t(10) = 100/1000 < 1$, so the TM could not even get started). This is reflected in the $n \geq c_M$ in the definition.

With this in mind, we now give the definition.

Definition 2.2. [Time complexity classes – boolean] Let $t : \mathbb{N} \rightarrow \mathbb{N}$ be a function. (\mathbb{N} denotes the natural numbers $\{0, 1, 2, \dots\}$.) $\text{TM-Time}(t)$ denotes the functions f that map bit strings x from a subset $X \subseteq [2]^*$ to a set $f(x)$ for which there exists a TM M such that, on any input $x \in X$ of length $\geq c_M$, M computes y within $t(|x|)$ steps and $y \in f(x)$.

$$P := \bigcup_{d \geq 1} \text{TM-Time}(n^d),$$

$$\text{Exp} := \bigcup_{d \geq 1} \text{TM-Time}(2^{n^d}).$$

We will not need to deal with relations and partial functions until later in this text.

Also, working with boolean functions, i.e., functions f with range $[2]$ slightly simplifies the exposition of a number of results we will see later. To avoid an explosion of complexity classes, we adopt the following convention.

Convention about complexity classes:

Unless specified otherwise, inclusions and separations among complexity classes refer to boolean functions. For example, an expression like $P \subseteq NP$ means that every boolean function in P is in NP .

As hinted before, the definition of P is robust. In the next few sections we discuss this robustness in more detail, and also introduce a number of other central computational models.

2.2 TMs with large alphabet

As our first example of robustness, we discuss TMs with arbitrary alphabet. This might seem like a detail, but in fact we are going to shortly present a cute problem in the area which will come up again and is, as far as I know, open. To set the stage, we first a relatively straightforward power-time simulation.

We define TMs *with alphabet size a* as in Definition 2.1 but with $|A|$ of size a ; we will only be interested in $a \geq |A|$ so we can think of adding symbols to A in Definition 2.1. We define similarly $\text{TM-Time}(t(n))$ with alphabet size a .

Theorem 2.1. $\text{TM-Time}(t(n))$ with alphabet size $a \subseteq \text{TM-Time}(c_a t(n) + c_a n^2)$.

Proof. Given a machine M_a as in the LHS, we construct machine M as in the RHS as follows. We use $c \log a \leq c_a$ tape symbols of M to encode one tape symbol of M_a . First we need to re-encode the input x . This takes time $c_a n^2$ as follows. First we move the head to the rightmost symbol of x in position $|x|$, and we shift it right of c_a positions. Then we go to the adjacent symbol in position $|x| - 1$, and shift all the contents to the right of this by c_a positions, to the right. We continue in this way.

Once this re-encoding is done, M can simulate M_a step-by-step, spending time c_a for each step of M_a . **QED**

This shows that the definition of P is robust w.r.t. different alphabet sizes. Yet the simulation is unsatisfactory due to the need of re-encode the input which gives a quadratic time blow-up.

Question 2.1. *Is the n^2 term in Theorem 2.1 necessary?*

2.2.1 The universal TM

Universal machines can simulate any other machine on any input. These machines play a critical role in some results we will see later. They also have historical significance: before them machines were tailored to specific tasks. One can think of such machines as epitomizing the victory of *software* over *hardware*: A single machine (hardware) can be programmed (software) to simulate any other machine.

Lemma 2.1. There is a TM U that on input (M, t, x) where M is a TM, t is an integer, and x is a string:

- Stops in time $|M|^c \cdot t \cdot |x|$,
- Outputs $M(x)$ if the latter stops within t steps on input x .

Proof. We maintain the invariant that M and t are always next to the tape head of U . After the simulation of each step of M the tape of U will contain

$$(x, M, \dot{i}, t', y)$$

where M is in state i , the tape of M contains xy and the head is on the left-most symbol of y . The integer t' is the counter decreased at every step. Computing the transition of M takes time $|M|^c$. Decreasing the counter takes time $c|t|$. To move M and t next to the tape head takes $c|M||t|$ time. **QED**

2.3 Multi-tape machines (MTMs)

Definition 2.3. A k -TM is like a TM but with k tapes, where the heads on the tapes move independently. The input is placed on the first tape, and all other tapes are initialized to $_$. The output is on the first tape. k -TM-Time is defined analogously to TM-Time. We write MTM for multi-tape machine for some number k of tapes.

Exercise 2.4. Prove that Palindromes is in 2-TM-Time(cn). Compare this to the run-time from the the TM in Exercise 2.1.

The following result implies in particular that P is unchanged if we define it in terms of TMs or k -TMs.

Theorem 2.2. k -TM-Time($t(n)$) \subseteq TM-Time($c_k t^2(n)$) for any $t(n) \geq n$ and k .

Exercise 2.5. Prove this. Recall that we defined TMs with fixed alphabet, and cf. section §2.2.

A much less obvious simulation is given by the following fundamental result about MTMs. It shows how to reduce the number of tapes to *two*, at little cost in time. Moreover, the head movements of the simulator are restricted in a sense that at first sight appears too strong.

Theorem 2.3. [92, 153] k -TM-Time($t(n)$) \subseteq 2-TM-Time($c_k t(n) \log t(n)$), for every function $t(n) \geq n$. Moreover, the 2-TM is *oblivious*: the movement of each tape head depends only on the length of the input.

Using this results one can prove the existence of universal MTMs similar to the universal TMs in Lemma 2.1. However, we won't need this result so we omit the proof.

2.4 Circuits

We now define circuits. It may be helpful to refer to figure 2.2 and figure 2.3.




Figure 2.2: A circuit computing the Xor of two bits.




Figure 2.3: An alternating circuit.

Definition 2.4. A *circuit*, abbreviated Ckt, is a directed acyclic graph where each node is one of the following types: an input variable (fan-in 0), an output variable (fan-in 1), a negation gate \neg (fan-in 1), an And gate \wedge (fan-in 2), or an Or gate \vee (fan-in 2). The *fan-in* of a gate is the number of edges pointing to the gate, the *fan-out* is the number of edges pointing away from the gate.

An *alternating circuit*, abbreviated AltCkt, or AC, is a circuit with unbounded fan-in Or and And gates arranged in alternating layers (that is, the gates at a fixed distance from the input all have the same type). For each input variable x_i the circuit has both x_i and $\neg x_i$ as input.

A DNF (resp. CNF) is an AltCkt whose output is Or (resp. And). The non-output gates are called *terms* (resp. *clauses*).

$\text{CktGates}(g(n))$ denotes the set of function $f : [2]^* \rightarrow [2]^*$ that, for all sufficiently large n , on inputs of length n have circuits with $g(n)$ gates; input and output gates are not counted. The *Size* of a circuit is the number of gates. We also define

$$\text{PCkt} := \bigcup_d \text{CktGates}(n^d).$$

We also denote by AC the class of functions computable by AC circuits of size n^d and depth d for a constant d .

Exercise 2.6. [Pushing negation gates at the input] Show that for any circuit $C : [2]^n \rightarrow [2]$ with g gates and depth d there is a monotone circuit C' (that is, a circuit without Not gates) with $2g$ gates and depth d such that for any $x \in [2]^n : C(x_1, x_2, \dots, x_n) = C'(x_1, \neg x_1, x_2, \neg x_2, \dots, x_n, \neg x_n)$.

Often we will consider computing functions on *small inputs*. In such cases, we can often forget about details and simply appeal to the following result, which gives exponential-size circuits which are however good enough if the input is really small. In a way, the usefulness of the result goes back to the locality of computation. The result, which is a circuit analogue of Exercise 2.3, will be extensively used in this book.

Theorem 2.4. Every function $f : [2]^n \rightarrow [2]$ can be computed by

- (1) circuits of size $\leq (1 + o(1))2^n/n$, and
- (2) A DNF or CNF with $\leq 2^n + 1$ gates (in particular, circuits of size $\leq n2^n$).

Exercise 2.7. Prove that the Or function on n bits has circuits of size cn . Prove Item (2) in Theorem 2.4. Prove a weaker version of Item (1) in Theorem 2.4 with bound $cn2^n$.

Exercise 2.8. Prove that the sum of two n -bit integers can be computed by circuits with cn gates, and by alternating circuits of depth c and size n^c .

We now show that circuits can simulate MTMs. We begin with a simple but instructive simulation of TMs which incurs a quadratic loss, then present a more interesting quasilinear simulation.

Theorem 2.5. Suppose an s -state TM computes $f : [2]^* \rightarrow [2]$ in time $t \geq n$. Then $f \in \text{CktGates}(c_s t^2(n))$. In particular

$$P \subseteq \text{PCkt}.$$

For this proof and the next it is convenient to represent a configuration of a TM in a slightly different way, as a string over the alphabet $A \times \{0, 1, \dots, s\}$. String

$$(a_1, 0)(a_2, 0) \dots (a_{i-1}, 0)(a_i, j)(a_{i+1}, 0) \dots (a_m, 0)$$

with $j > 0$ indicates that (1) the tape content is $a_1 a_2 \dots a_m$ with blanks on either side, (2) the machine is in state \underline{j} , and (3) the head of the machine is on the i tape symbol a_i in the string.

Locality of computation here means that one symbol in a string only depends on the symbols corresponding to the same tape cell i in the previous step and its two neighbors – three symbols total – because the head only moves in one step.

Proof of Theorem 2.5. Given a TM M with s states consider a $t \times (2t + 1)$ matrix T , a.k.a. the *computation table*, where row i is the configuration at time i . The starting configuration in Row 1 has the head in the middle cell. Note we don't need more than t cells to the right or left because the head moves only by one cell in one step. Next we claim that Row $i + 1$ can be computed from Row i by a Ckt with $c_s t$ gates. This follows by locality of computation, where note each entry in Row $i + 1$ can be computed by a Ckt of size c_s , by Theorem 2.4.

Stacking t such circuits we obtain a circuit of size $c_s t^2$ which computes the end configuration of the TM.

There remains to output the value of the function. Had we assumed that the TM writes the output in a specific cell, we could just read it off by a circuit of size c . Without the assumption, we can have a circuit $C : A \times \{0, 1, \dots, s\} \rightarrow [2]$ which outputs 1 on (x, y) iff $y \neq 0$ and $x = 1$ (i.e., if x is a 1 that is under the TM's head). Taking an Or such circuits applied to every entry in the last row of T concludes the proof. **QED**

The simulation in Theorem 2.5. incurs a quadratic loss. However, a better simulation exists. In fact, this applies even to k -TMs.

Theorem 2.6. Suppose an s -state k -TM computes $f : [2]^* \rightarrow [2]$ in time $t(n) \geq n$. Then $f \in \text{CktGates}(c_{s,k} t(n) \log t(n))$.

Exercise 2.9. Prove Theorem 2.6 assuming Theorem 2.3.

Next we give a direct proof that doesn't need Theorem 2.3.

Proof. We prove this for $k = 1$, the extension to larger k does not need new ideas and is omitted. Given a TM M , we construct a circuit S_m that on input a configuration of M with m tape symbols where the head position is within $m/4$ symbols from the center, it computes the configuration reached by M after $m/4$ steps of the computation.

We shall give an inductive construction of S_m satisfying

$$\text{Size}(S_m) \leq 2 \cdot \text{Size}(S_{m/2}) + cm$$

with base case $\text{Size}(S_c) \leq c_M$. This implies $\text{Size}(S_t) \leq c_M t \log t$, as desired.

To construct S_m we think of the m symbols as divided into c blocks, and we rely on a couple of auxiliary circuits. Circuit H_m given an m -symbol configuration computes in which block the head is; circuit R_m given an m -symbol configuration and $i \leq c$, rotates the blocks by i positions. We can now program S_m as follows. First run H_m to get in which block i the head is. Use R_m to rotate the blocks by i positions so that the head is in a block closest to the middle. Run $S_{m/2}$. Now again run H_m to get j , and then R_m to move block j closest to the middle. Run $S_{m/2}$. Finally, use R_m to restore the blocks by rotating them back by $i + j$ positions.

This circuit simulates $(m/2)/4 + (m/2)/4 = m/4$ steps, as desired. The circuits R and H can be implemented using cm gates. **QED**

In the other direction, TMs can simulate circuits if they have enough states. In general, allowing for the number of states to grow with the input length gives models with “hybrid uniformity.”

Exercise 2.10. Suppose that $f : [2]^n \rightarrow [2]$ has circuits with s gates. Show that f can be computed by a TM with s^c states in time s^c .

2.5 Rapid-access machines (RAMs)

“In some sense we are therefore merely making concrete intuitions that already pervade the literature. A related model has, indeed, been treated explicitly” [...] [14]

The main feature that’s missing in all models considered so far is the ability to read and write a memory cell in one time step given the address. This feature is called *direct addressing*, and is common place in programming languages (where for example we define an array A and then we can access cell i in the array via $A[i]$).

One can augment TMs with this capability by equipping them with an extra *addressing tape* and a special “jump” state which causes the head on a tape to move in one step to the address on the address tape. This model is simple enough to define, and could in a philosophical sense be the right model for how hardware can scale, since we charge for the time to write the address. However, other models are closer to how computers seem to operate, at least over small inputs. We want to think of manipulating small integers and addresses as constant-time operations, as one typically has in mind when programming. There is a variety of such models, and some arbitrariness in their definition. Basically, we want to think of the memory as an array μ of s cells of w bits and allow for typical operations of them, including addressing arithmetic and *indirect addressing*: reading and writing the cell indexed by another cell.

One issue that arises is how much memory the machine should have and consequently how big w should be. There are two main options here. For “typical programming,” we have a fixed memory size s and time bound t in mind, for example $s = n^3$ and $t = n^2$. A good choice then is to set $w := \lceil \log_2 s \rceil$ bits.

This however makes it harder to compare machines with different memory bounds. Also in some scenarios the memory size and the time bound are not fixed. This occurs for example when simulating another machine. To handle such scenarios we can start with a memory of $s = n + c$ cells, and a cell size of $w = \lceil \log_2 s \rceil$ bits, enough to access the input. We then equip machines with the operation MAlloc which increases the memory (i.e., s) by one, and always sets $w := \lceil \log_2 s \rceil$. Note the latter operation may increase w by 1. The MAlloc operation is akin to the TM’s tape head wandering into unknown cells.

There are also two options for how the input is given to the machine. The difference doesn’t matter if you don’t care about w factors in time, but it matters if you do. For many problems, like sorting, etc. we think of the input and the output as coming in n cells of w bits. (Typically, $w = c \log n$, and one can simulate such cells with c cells with $\log n$ bits.) In this case, the RAM is computing a function $f : ([2]^w)^n \rightarrow ([2]^w)^m$ and the input to the RAM is given accordingly. This is what one often has in mind when writing programs that involve numbers. For other problems, it is natural to just give one bit of the input in each cell. That is, the RAM is computing $f : [2]^n \rightarrow [2]^m$ and bit i of the input is placed in the i input cells. We will not be concerned too much with small factors and so we pick the second choice for simplicity. This choice will also make it easier later to write computation in certain useful formats (cf. Lemma 6.1).

Definition 2.5. A w -bit ℓ -line rapid-access machine (RAM) with s cells consists of a memory array $\mu[1..s]$ of s cells of w bits, c registers r_1, r_2, \dots of w bits, and a program of ℓ lines.

Each line of the program contains an instruction among the following:

- Standard arithmetical, logical, and control-flow operations, such as $r_1 = r_2 + r_3$, if $r_1 = 0$ then goto line 17, etc.
- $r_i := \mu[r_j]$, called a Read operation, which reads the r_j memory cell and copies its content into r_i ,
- $\mu[r_i] := r_j$, called a Write operation, which writes r_j into memory cells r_i , memory cell and copies its content into r_i ,
- MAlloc which increases s by 1 and, if $s \geq 2^w$ also increases w by 1,
- Stop.

Read and write operations out of boundary indices have no effect.

On an input $x \in [2]^n$, the RAM starts the computation with $s := n + 1$ cells of memory. The input is written in cells $1..n$, while $\mu[0]$ contains the length n of the input.

The output is written starting in cell 1.

We use RAMs as our main model for time inside P.

Definition 2.6. $\text{Time}(t(n))$ is defined as $\text{TM-Time}(t(n))$ but for RAMs instead of TMs.

Theorem 2.7. $\text{Time}(t(n)) \subseteq \text{TM-Time}(t^c(n))$, for any $t(n) \geq n$.

Exercise 2.11. Prove it.

What is the relationship between circuits and RAMs? If a “description” of the circuit is given, then a RAM can simulate the circuit efficiently. The other way around is not clear. It appears that circuits need a quadratic blow-up to simulate RAMs.

Exercise 2.12. Give a function $f : [2]^* \rightarrow [2]$ in $\text{Time}(c \log n)$ but which requires circuits of size $\geq cn$.

There are universal RAMs that can simulate any other RAM with only a constant-factor overhead, unlike the logarithmic-factor overhead for tape machines.

Lemma 2.2. There is a RAM U that on input (P, t, x) where P is a RAM, t is an integer, and x is an input

- Stops in time ct ,
- Outputs $P(x)$ if the latter stops within t steps on input x .

Proof. Throughout the computation, U will keep track of the memory size s_P and cell-size w_P of P . These are initialized as in the initial configuration of P on input x , whereas U starts with bigger values, since its input also contains P and t . Let h be the first cell where the input x starts. Memory location i of P is mapped to $i+h$ during the simulation. When P performs an operations among registers, U simulates that with its own registers, but discards the data that does not fit into w_P bits.

After each step, U decreases the counter. The counter can be stored in t cells, one bit per cell. The total number of operations to decrease such a counter from t to 0 is $\leq ct$. Alternatively, we can think of the counter as being stored in a single register at the beginning of the simulation. Then decreasing the counter is a single operation. **QED**

2.6 A challenge to the computability thesis

Today, there’s a significant challenge to the computability thesis. This challenge comes from... I know what you are thinking: *Quantum computing, superposition, factoring*. Nope. *Randomness*.

The last century or so has seen an explosion of randomness affecting much of science, and computing has been a leader in the revolution. Today, randomness permeates computation. Except for basic “core” tasks, using randomness in algorithms is standard. So let us augment our model with randomness.

Definition 2.7. A *randomized* (or *probabilistic*) RAM, written RRAM, is a RAM equipped with the extra instruction

- $r_i := \text{Rand}$, which sets r_i to a uniform value, independent of all previous random choices.

For a RRAM and a sequence $R = R_1, R_2, \dots$ we write $M(x, R)$ for the execution of M on input x where the j -th instruction $r_i := \text{Rand}$ is replaced with $r_i := R_j$.

We refer to $\text{BPTIME}(t(n))$ with error $\epsilon(n)$ as the set of functions f that map bit strings x from a subset $X \subseteq [2]^*$ to a set $f(x)$ for which there exists a RRAM M such that, on any input $x \in X$ of length $\geq c_M$, M stops within $t(|x|)$ steps and $\mathbb{P}_R[M(x, R) \in f(x)] \geq 1 - \epsilon(|x|)$.

If the error ϵ is not specified then it is assumed to be $1/3$. Finally, we define

$$\text{BPP} := \bigcup_a \text{BPTIME}(n^a).$$

Exercise 2.13. Does the following algorithm show that deciding if a given integer x is prime is in BPP? “Pick a uniform integer $y \in [2..x - 1]$. If y divides x return NOT PRIME, else return PRIME.”

Today, one usually takes BPP, not P, for “feasible computation.” Thus it is natural to investigate how robust BPP is.

2.6.1 Robustness of BPP: Error reduction and tail bounds for the sum of random variables

The error in the definition of BPTIME is somewhat arbitrary because it can be reduced. The way you do this is natural. For boolean functions, you repeat the algorithm many times, and take a majority vote. To analyze this you need probability bounds for the deviation of the sum of random variables (corresponding to the outcomes of the algorithm) from the mean. Such deviation bounds permeate theoretical computer science, and many other fields as well.

Theorem 2.8. Let X_1, X_2, \dots, X_t be i.i.d. boolean random variables with $p := \mathbb{P}[X_i = 1]$. Then for $q \geq p$ we have $\mathbb{P}[\sum_{i=1}^t X_i \geq qt] \leq 2^{-D(q|p)t}$, where

$$D(q|p) := q \log \left(\frac{q}{p} \right) + (1 - q) \log \left(\frac{1 - q}{1 - p} \right)$$

is the *divergence*.

The proof uses the following basic facts.

Claim 2.1. Let X be a real-valued r.v. s.t. $X \geq 0$ always. Then $\mathbb{P}[X \geq t] \leq \mathbb{E}[X]/t$ for every $t > 0$.

Exercise 2.14. Prove this. Hint: Use that for any event E , $\mathbb{E}[X] = \mathbb{E}[X|E]\mathbb{P}[E] + \mathbb{E}[X|\text{not } E]\mathbb{P}[\text{not } E]$.

Claim 2.2. If X and Y are independent, real-valued random variables then $\mathbb{E}[X \cdot Y] = \mathbb{E}[X] \cdot \mathbb{E}[Y]$.

Proof of Theorem 2.8. For $z \geq 1$, to be picked later, the function $x \rightarrow z^x$ is increasing. Using this and then Claim 2.1 and finally the independence of the X_i , the LHS equals

$$\mathbb{P}[z^{\sum_{i=1}^t X_i} \geq z^{qt}] \leq \frac{\mathbb{E}[z^{\sum_{i=1}^t X_i}]}{z^{qt}} = \frac{\prod_{i=1}^t \mathbb{E}[z^{X_i}]}{z^{qt}} = \left(\frac{pz + 1 - p}{z^q} \right)^t =: b^t.$$

To minimize b we set

$$z := \frac{q(1-p)}{(1-q)p},$$

which is ≥ 1 because $q \geq p$, and obtain

$$b = \frac{1-p}{z^q} = \left(\frac{p}{q} \right)^q \left(\frac{1-p}{1-q} \right)^{1-q}.$$

QED

Now one can get a variety of bounds by bounding divergence for different settings of parameter. We state one such bound which we use shortly.

Fact 2.1. $D(q|p) \geq c(p-q)^2$, for any $p, q \in [0, 1]$.

Exercise 2.15. For $q = 1/2$ and $p = 1/2 - \epsilon$ plot both sides of Fact 2.1 as a function of ϵ . (Hint: I used <https://www.desmos.com/calculator>)

The proof of the tail-bound Theorem 2.8 is flexible and applies to a variety of useful settings. The most interesting extensions concern *dependent* random variables, where in general the bounds are weaker. In the next exercise we instead consider other settings where the bounds in Theorem 2.8 continue to hold.

Exercise 2.16. Prove that the tail bound in Theorem 2.8 holds as stated more generally for any independent random variables X_1, X_2, \dots, X_t distributed in $[0, 1]$ with $p := \sum_i \mathbb{E}[X_i]/t$. Guideline: Repeat the same proof as before. Use that $z^x \leq 1 + x(z-1)$ and the arithmetic-mean geometric mean inequality (AM-GM) inequality: for all $a_i \geq 0$: $(\sum_{i \in [t]} a_i)/t \geq (\prod_{i \in [t]} a_i)^{1/t}$.

Now suppose the X_i are more generally distributed in $[a, b]$. For $q = \epsilon + p$ prove a deviation bound of $2^{-c\epsilon^2 t / (b-a)^2}$.

Using the tail-bound Theorem 2.8 we can prove the error reduction stated earlier.

Theorem 2.9. [Error reduction for BPP] For boolean functions, the definition of BPP (Definition 2.7) remains the same if $1/3$ is replaced with $1/2 - 1/n^a$ or $1/2^{n^a}$, for any constant a .

Proof. Suppose that f is in BPP with error $p := 1/2 - 1/n^a$ and let M be the corresponding RRAM. On an input x , let us run $t := n^{2a} \cdot n^b$ times M , each time with fresh randomness, and take a majority vote. The new algorithm is thus

$$\text{Maj}(M(x, R_1), M(x, R_2), \dots, M(x, R_t)).$$

This new algorithm makes a mistake iff at least $t/2$ runs of M make a mistake. To analyze this error probability we invoke Theorem 2.8 where $X_i := 1$ iff run i of the algorithm makes a mistake, i.e., $M(x, R_i) \neq f(x)$, and $\epsilon := 1/n^a$. By Fact 2.1 we obtain an error bound of

$$2^{-D(1/2|1/2-\epsilon)t} \leq 2^{-\epsilon^2 t} \leq 2^{-n^b},$$

as desired. The new algorithm still runs in power time, for fixed a and b . **QED**

Exercise 2.17. Consider an alternative definition of BPTIME, denoted BPTIME', which is analogous to BPTIME except that the requirement that the machine always stops within $t(|x|)$ steps is relaxed to “the *expected* running time of the machine is $t(|x|)$.”

Show that defining BPP with respect to BPTIME or BPTIME' is equivalent.

Exercise 2.18. Consider *biased* RRAMs which are like RRAMs except that the operation Rand returns one bit which, independently from all previous calls to Rand, is 1 with probability $1/3$ and 0 with probability $2/3$. Show that BPP does not change if we use biased RRAMs.

2.6.2 Does randomness buy time?

We can always brute-force the random choices in exponential time. If a randomized machine uses r random bits then we can simulate it deterministically by running it on each of the 2^r choices for the bits. A RRAM machine running in time $t \geq n$ has registers of $\leq c \log t$ bits. Each Rand operation gives a uniform register, so the machine uses $\leq ct \log t$ bits. This gives the following inclusions.

Theorem 2.10. $\text{Time}(t) \subseteq \text{BPTIME}(t) \subseteq \text{Time}(c^{t \log t})$, for any function $t = t(n)$. In particular, $\text{P} \subseteq \text{BPP} \subseteq \text{EXP}$.

Proof. The first inclusion is by definition. The idea for the second was discussed before, but we need to address the detail that we don't know what t is. One way to carry through the simulation is as follows. The deterministic machine initializes a counter r to 0. For each value of r it enumerates over the 2^r choices R for the random bits, and runs the RRAM on

each choice of R , keeping track of its output on each choice, and outputting the majority vote. If it ever runs out of random bits, it increases r by 1 and restarts the process.

To analyze the running time, recall we only need $r \leq ct \log t$. So the simulation runs the RRAM at most $ct \log t \cdot 2^{ct \log t} \leq 2^{ct \log t}$ times, and each run takes time ct , where this last bound takes into account the overhead for incrementing the choice of r , and redirecting the calls to Rand to R . **QED**

Now, two surprises. First, $\text{BPP} \subseteq \text{EXP}$ is the fastest deterministic simulation we can *prove* for RAMs, or even 2-TMs. On the other hand, what may come as a bigger surprise, it appears that many people *believe* that in fact $\text{P} = \text{BPP}$! Moreover, it appears commonly believed that the overhead to simulate randomized computation deterministically is very small. Here the mismatch between our ability and common belief is abysmal.

However, we can do better for TMs. A *randomized* TMs has two transition functions σ_0 and σ_1 , where each is as in Definition 2.1. At each step, the TM uses σ_0 or σ_1 with probability $1/2$ each, corresponding to tossing a coin. We can define TM-BPTime as BPTime but with randomized TMs instead of RRAMS.

Theorem 2.11. $\text{TM-BPTime}(t) \subseteq \text{Time}(2^{\sqrt{t} \log^c t})$, for any $t = t(n) \geq n$.

2.6.3 Polynomial identity testing

We now discuss an important problem which is in BPP but not known to be in P. In fact, in a sense to be made precise later, this is *the* problem in BPP which is not known to be in P. To present this problem we introduce two key concepts which will be used many times: *finite fields*, and *arithmetic circuits*.

Finite fields A finite field \mathbb{F} is a finite set with elements 0 and 1 that is equipped with operations $+$ and \cdot that behave “in the same way” as the corresponding operations over the reals \mathbb{R} or the rationals \mathbb{Q} , which are *infinite* fields. One example are the integers modulo a prime p . For $p = 2$ this gives the field with two elements where $+$ is Xor and \cdot is And. For larger p you add and multiply as over the integers but then you take the result modulo p .

The following summarizes key facts about finite fields. The case of prime fields suffices for the main points of this section, but stating things for general finite fields actually simplifies the exposition overall (since otherwise we need to add qualifiers to the size of the field).

Fact 2.2. [Finite fields] A unique finite field of size q exists iff $q = p^t$ where p is a prime and $t \in \mathbb{N}$. This field is denoted \mathbb{F}_q .

Elements in the field can be identified with $\{0, 1, \dots, p - 1\}^t$.

Given q , one can compute a *representation* of a finite field of size q in time $(tp)^c$. This representation can be identified with p plus an element of $\{0, 1, \dots, p - 1\}^t$.

Given a representation r and field elements x, y computing $x+y$ and $x \cdot y$ is in $\text{Time}(n \log^c n)$.

Fields of size 2^t are of natural interest in computer science. It is often desirable to have very explicit representations for such and other fields. Such representations are known and are given by simple formulas, and are in particular computable in linear time.

Example 2.3. We can represent the elements of \mathbb{F}_{p^t} as (the coefficients of) polynomials of degree $< t$ over \mathbb{F}_p . Addition is done component-wise, and multiplication occurs modulo an irreducible polynomial of degree t over the base field \mathbb{F}_p , i.e., a polynomial that cannot be factored as the product of two non-constant polynomials. It is known $z^t + z^{t/2} + 1$ is irreducible over \mathbb{F}_2 for any $t = 2 \cdot 3^\ell$ for any ℓ , giving very explicit representations. For example, consider the field elements $z^2 + 1$ and $z^{t-1} + 1$ over such a representation of \mathbb{F}_{2^t} . Their sum equals $z^{t-1} + z^2$, and their product equals $z^{t+1} + z^2 + z^{t-1} + 1 = z^{t-1} + z^{t/2+1} + z^2 + z + 1$.

Arithmetic circuits We now move to defining arithmetic circuits, which are a natural generalization of the circuits we encountered in section §2.4.

Definition 2.8. An *arithmetic circuit* over a field \mathbb{F} is a circuit where the gates compute the operations $+$ and \cdot over \mathbb{F} , or are constants, or are input variables. Such a circuit computes a polynomial mapping $\mathbb{F}^n \rightarrow \mathbb{F}$.

The PIT (polynomial identity testing) problem over \mathbb{F} : Given an arithmetic circuit C over \mathbb{F} with n input variables, does $C(x) = 0$ for every $x \in \mathbb{F}^n$?

The PIT problem *over large fields* is in BPP but it is not known to be in P. The requirement that the field be large is critical, see Problem 4.1.

Theorem 2.12. [PIT over large fields in BPP] Given an arithmetic circuit C and the representation of a finite field of size $\geq c2^{|C|}$ we can solve PIT in BPP.

To prove this theorem we need the following fundamental fact.

Lemma 2.3. Let p be a polynomial over a field \mathbb{F} with n variables and degree $\leq d$. Let S be a finite subset of \mathbb{F} , and suppose $d < |S|$. The following are equivalent:

1. p is the zero polynomial.
2. $p(x) = 0$ for every $x \in \mathbb{F}^n$.
3. $\mathbb{P}_{x_1, x_2, \dots, x_n \in S}[p(x) = 0] > d/|S|$.

Proof of Lemma 2.3. The implications 1. \Rightarrow 2. \Rightarrow 3. are trivial, but note that for the latter we need $d < |S|$. The implication 3. \Rightarrow 1. is not trivial. We proceed by induction on n .

The base case $n = 1$ is the fact that if p has more than d roots then it is the zero polynomial. This fact in turn can be proved by induction on the degree. The base case $d = 0$ is obvious. For larger d , suppose a is a root of p and use division for polynomials to write $p = (x - a)q + r$ where q has degree $\leq d - 1$ and $r \in \mathbb{F}$. Because a is a root we have $r = 0$, and so $p = (x - a)q$ and q has $d - 1$ roots, and by induction $q = 0$ and so $p = 0$.

For larger n write

$$p(x_1, x_2, \dots, x_n) = \sum_{i=0}^d x_1^i p_i(x_2, x_3, \dots, x_n).$$

If p is not the zero polynomial then there is at least one i such that p_i is not the zero polynomial. Let j be the largest such i . Note that p_j has degree at most $d - j$. By induction hypothesis

$$\mathbb{P}_{x_2, \dots, x_n \in S}[p_j(x) = 0] \leq (d - j)/|S|.$$

For every choice of x_2, x_3, \dots, x_n s.t. $p_j(x) \neq 0$, the polynomial p is a non-zero polynomial $q_{x_2, x_3, \dots, x_n}(x_1)$ only in the variable x_1 . Moreover, its degree is at most j by our choice of j . Hence by the $n = 1$ case the probability that q is 0 over the choice of x_1 is $\leq j/|S|$.

Overall,

$$\mathbb{P}_{x_1, x_2, \dots, x_n \in S}[p(x) = 0] \leq (d - j)/|S| + j/|S| = d/|S|.$$

QED

Exercise 2.19. Show that the equivalence between 1. and 2. does not hold over small fields such as \mathbb{F}_2 and large d .

Proof of Theorem 2.12. A circuit C contains at most $|C|$ multiplication gates. Each multiplication gate at most squares the degree of its inputs. Hence C computes a polynomial of degree $\leq 2^{|C|}$. Let S be a subset of size $c \cdot 2^{|C|}$ of \mathbb{F} . Assign uniform values from S independently to each variables, and evaluate the circuit. If C evaluates to 0 everywhere then obviously the output will be 0. Otherwise, by Lemma 2.3, the probability we get a 0 is $\leq 2^{|C|}/c2^{|C|} \leq 1/3$. **QED**

To show that the PIT problem over the *integers* is in BPP the following result is useful.

Theorem 2.13. [Prime number theorem] $\lim_{n \rightarrow \infty} (\text{Number of primes} \leq n) / (n / \log_e n) = 1$.

As is often the case in computer science, we don't need the full strength of Theorem 2.13. An approximate version with $\log_e n$ replaced by $\log^c n$ suffices, and it has a considerably easier proof. (Moreover sometimes easier proofs are easier to adapt to other settings of interest in computer science.) This weak version is stated next.

Theorem 2.14. [Weak prime number theorem] The number of primes in $[t]$ is $\geq t / \log^c t$, for every $t \geq c$.

Exercise 2.20. Show that the PIT problem over the *integers* is in BPP. (Hint: Use Theorem 2.13 and that checking if a number is prime is in P.)

2.6.4 Simulating BPP by circuits

While we don't know if $P = BPP$, we can prove that, like P, BPP has power-size circuits.

Theorem 2.15. $BPP \subseteq PCKt$.

Proof. Let $f : X \subseteq [2]^* \rightarrow [2]$ be in BPP. By Theorem 2.9 we can assume that the error is $\epsilon < 2^{-n}$, and let M be the corresponding RRAM. Note

$$\mathbb{P}_R[\exists x \in [2]^n : M(x, R) \neq f(x)] \leq \sum_{x \in [2]^n} \mathbb{P}_R[M(x, R) \neq f(x)] \leq 2^n \cdot \epsilon < 1,$$

where the first inequality is a union bound.

Therefore, there is a fixed choice for R that gives the correct answer for every input $x \in [2]^n$. This choice can be hardwired in the circuit, and the rest of the computation can be written as a circuit by Theorem 2.5. **QED**

Exercise 2.21. In this exercise you will practice the powerful technique of combining tail bounds with union bounds, which was used in the proof of Theorem 2.15, and also see a related application of Lemma 2.3 .

An *error-correcting code* with block length n , message length k , minimum distance d , over the alphabet q , written $(n, k, d)_q$ is a subset $C \subseteq [q]^n$ of size q^k s.t. for any distinct $x, y \in C$, x and y differ in at least d coordinates.

(1) Prove the existence of $(n, an, bn)_2$ codes, for some constants $a, b > 0$ and every n .

(2) Given a prime power q , and $k \leq q$ construct an explicit $(q, k, q - k)_q$ code using Lemma 2.3. For explicitness, show that given q and $x \in [q]^k$ computing the corresponding codeword is in P.

2.6.5 Questions raised by randomness

The introduction of randomness in our model raises several fascinating questions. First, does “perfect” randomness exist “in nature?” Second, do we need “perfect” randomness for computation? A large body of research has been devoted to greatly generalize Problem 2.18 to show that, in fact, even imperfect sources of randomness suffices for computation. Third, do we need randomness at all? Is $P = BPP$?

One of the exciting developments of complexity theory has been the connection between the latter question and the “grand challenge” from the next chapter. At a high level, it has been shown that explicit functions that are hard for circuits can be used to *de-randomize* computation. In a nutshell, the idea is that if a function is hard to compute then its output is “random,” so can be used instead of true randomness. The harder the function the less randomness we need. At one extreme, we have the following striking connection:

Theorem 2.16. Suppose for some $a > 0$ there is a function in $\text{Time}(2^{an})$ which on inputs of length n cannot be computed by circuits with $2^{n/a}$ gates, for all large enough n . Then $P = BPP$.

In other words, either randomness is useless for power-time computation, or else circuits can speed up exponential-time uniform computation!

We will prove this in Chapter 13, Exercise 13.17.

2.7 Inclusion extend “upwards,” separations downwards

To develop intuition about complexity, we now discuss a general technique known as *padding*. In short, the technique shows that if you can trade resource X for Y , then you can also trade *a lot of* X for *a lot of* Y . For a metaphor, if you have a magical device that can turn one pound of sill into gold, you can also use it to turn *two* pounds of sill into gold. The contrapositive is that if you *can't* trade a lot of X for a lot of Y , then you also can't trade a little of X for a little of Y .

We give a first example using the classes that we have encountered so far.

Example 2.4. Suppose that $\text{BPTIME}(cn) \subseteq \text{TIME}(n^2)$. Then $\text{BPTIME}(n^2) \subseteq \text{TIME}(cn^4)$.

Proof. Let $f : [2]^* \rightarrow [2]$ be a function in $\text{BPTIME}(n^2)$. Consider the function f' that on input x of length n equals f computed on the first \sqrt{n} bits of x . Thus, inputs to f' are padded with $n - \sqrt{n}$ useless symbols.

Note that $f' \in \text{BPTIME}(cn)$, since in linear time we can erase the last $n - \sqrt{n}$ symbols and then just run the algorithm for f which takes time quadratic in \sqrt{n} which is linear in n . (If computing square roots is not an available instruction, one can show that computing \sqrt{n} can be done in linear time, for example using binary search.)

By assumption, $f' \in \text{TIME}(n^2)$.

To compute f in time cn^4 we can then do the following. Given input x of length n , pad x to an input of length n^2 in time cn^2 . Then run the algorithm for f' . This will take time $c(n^2)^2 \leq cn^4$. **QED**

2.8 Problems

Problem 2.1. [Indexing] Describe a TM that on input $(x, i) \in [2]^n \times \{1, 2, \dots, n\}$ outputs bit i of x in time $cn \log n$.

Problem 2.2. [Indexing] Describe a circuit with cn gates that on input $(x, i) \in [2]^n \times \{1, 2, \dots, n\}$ outputs bit i of x .

Problem 2.3. Show that Palindromes can be solved in time $n \log^c n$ on a randomized TM. (Yes, only one tape.)

Hint: View the input as coefficients of polynomials.

Problem 2.4. Give a function $f : X \subseteq [2]^* \rightarrow [2]$ that is in $\text{BPTIME}(c)$ but not in $\text{TIME}(n/100)$.

Problem 2.5. Let $f : [2]^n \rightarrow [2]$ be computable by an s -state k -TM in time t . Think of the input as m cells of w bits, so that $n = mw$. Consider circuits made of arbitrary functions which take as input $c_{s,k}$ cells and output one cell. (Each wire in this circuit carries one cell – the bits cannot be “broken up,” but the result would be non-trivial even if they could.) Show that f can be computed by such circuits of size $(t/w)^c$. For example, if $t = 100n$ and $w = 0.01n$ we have circuits of a constant number of gates.

Problem 2.6. For a circuit C on n bits denote by p_C the probability $\mathbb{P}_x[C(x) = 1]$.

(1) Show how to efficiently approximate p_C . Specifically: Give a power-time randomized algorithm that on input a circuit C and $\epsilon > 0$ written in unary (for example, as a string of $1/\epsilon$ ones) outputs p s.t. $|p - p_C| \leq \epsilon$ w.p. ≥ 0.9 .

(2) Show that the following decision version of (1) is in BPP: Given a circuit C , a number p (written in binary), and $\epsilon > 0$ written in unary, such that $|p_C - p| \geq \epsilon$, decide if $p_C \geq p$.

(3) What happens if in (2) you replace the assumption that $|p_C - p| \geq \epsilon$ with $|p_C - p| > 0$?

(4) Assume $P = BPP$ (recall this only refers to boolean functions, see the convention on page 18). Show how the approximation in (1) can be computed in P .

2.9 Notes

“It’s all over.”

The fundamental work on complexity is [65]. That work formalized computation for the first time, and discovered its self-referential ability, essentially inventing universal machines and the diagonalization technique (cf. section §3.3). Of course, [65] did not come out of nowhere, but was in fact a reaction to a program of automating mathematics, and it built on logical formalizations of mathematics; and diagonalization has its roots in (and takes the name from) the proof that the real numbers are uncountable. Also, there are several previous works aimed at formalizing computation in various branches of science. See [136] for an account of this compelling history. Still, if a fundamental work must be picked, [65] seems appropriate, for it can be considered the first work on impossibility results about general computation.

The formalization of computation in [65] is in terms of recursive functions, not unlike modern functional programming languages. As we saw, many other equivalent formalizations came about later. Tape machines were introduced in [193] and are closer to computer hardware or imperative programming languages. They also make it a little more intuitive to measure time and space in computation.

The brute-force computation of functions via circuits, Theorem 2.4, goes back to [172], see also [129].

—

For more on the circuit model in Problem 2.5 see [81].

For a computer-science friendly exposition of deviation bounds see the book [56].

—

Theorem 2.6 is towards the end of [153].

Theorem 2.11 is from [216].

Theorem 2.15 is from [3].


Theorem 2.16 is from [102].

-

The reference for Fact 2.2 is [174]. For more on finite fields see [122], for the fields \mathbb{F}_{2^t} in Example 2.3 see Theorem 1.1.28 in [199].

Chapter 3

The grand challenge



As mentioned in Chapter 1, our ability to prove impossibility results related to efficient computation appears very limited. We can now express this situation more precisely with the models we've introduced since then.

It is consistent with our knowledge that any problem in a standard algorithm textbook can be solved

1. in Time cn^2 on a TM, and
2. in Time cn on a 2-TM, and
3. by circuits of size cn .

Note that 2. implies 1. by Theorem 2.2, and many other relationships have been explored in Chapter 2.

In this chapter we begin to present several impossibility results, covering a variety of techniques which will be used later as well. As hinted above, they appear somewhat weak. However, jumping ahead, there is a flip side to all of this:

1. At times, contrary to our intuition, stronger impossibility results are actually *false*. One example appears in Chapter 1. A list will be given later.

2. Many times, the impossibility results that we can prove turn out to be, surprisingly, just “short” of proving major results. Here by “major result” I mean a result that would be phenomenal and that was in focus long before the connection was established. We will see several examples of this (section §7.3, section §10.2).
3. Yet other times, one can identify broad classes of proof techniques, and argue that impossibility results can’t be proved with them (section §19.1).

Given this situation, I don’t subscribe to the general belief that stronger impossibility results are true and we just can’t prove them.

3.1 Information bottleneck: Palindromes requires quadratic time on TMs

Intuitively, the weakness of TMs is the bottleneck of passing information from one end of the tape to the other. We now show how to formalize this and use it show that deciding if a string is a palindrome requires *quadratic* time on TMs, which is tight and likely matches the time in Exercise 2.1. The same bound can be shown for other functions; palindromes just happen to be convenient to obtain matching bounds.

Theorem 3.1. Palindromes \notin TM-Time($t(n)$) for any $t(n) = o(n^2)$.

More precisely, for every n and s , an s -state TM that decides if an n -bit input is a palindrome requires time $\geq cn^2/\log s$.

The main concept that allows us to formalize the information bottleneck mentioned above is the following.

Definition 3.1. A *crossing sequence* of a TM M on input x and boundary i , abbreviated i -CS, is the sequence of states that M is transitioning to when crossing cell boundary i (i.e., going from Cell i to $i + 1$ or vice versa) during the computation on x .

Example 3.1. We think of a step of a TM as first changing state and then moving the head. We write $u\overset{i}{v}w$ if the tape content is uvw and the TM is in state i with the head on v , where $u, w \in A^*$ and $v \in A$, cf. Definition 2.1. The computation

$\overset{0}{0}$	0	0	0
#	$\overset{5}{0}$	0	0
#	0	$\overset{2}{0}$	0
#	0	x	$\overset{2}{0}$
#	0	$\overset{2}{x}$	0
#	$\overset{1}{0}$	x	0
#	0	$\overset{7}{x}$	0

has the 2-cs (marked with double vertical line; first column is cell 1) equal to 2, 1, 7.

The idea in the proof is very interesting. If M accepts inputs x and y and those two inputs have the same i -CS for some i , then we can “stitch together” the computation of M on x and y at boundary i to create a new input z that is still accepted by M . The input z is formed by picking bits from x to the left of cell boundary i and bits from y to the right of i :

$$z := x_1x_2 \cdots x_i y_{i+1} y_{i+2} \cdots y_n.$$

The proof that z is still accepted is left as an exercise.

Example 3.2. The following computation has the same 2-cs as the previous example

$\frac{0}{0}$	1	1	0
0	$\frac{4}{1}$	1	0
0	0	$\frac{2}{1}$	0
0	$\frac{1}{0}$	0	0
0	0	$\frac{7}{0}$	0

If $\underline{7}$ is the accept state, then the TM would also accept the “stitched” input

0010

because on that input the TM has the following “stitched” computation:

$\frac{0}{0}$	0	1	0
#	$\frac{5}{0}$	1	0
#	0	$\frac{2}{1}$	0
#	$\frac{1}{0}$	0	0
#	0	$\frac{7}{0}$	0

Note that the number of steps of the stitched computations needs not be the same.

Now, for many problems, stitched input z should *not* be accepted by M , and this gives a contradiction. In particular this will be the case for palindromes. We are going to find two palindromes x and y that have the same i -CS for some i , but the corresponding z is not a palindrome, yet it is still accepted by M . We can find these two palindromes if M takes too little time. The basic idea is that if M runs in time t , because i -CSs for different i correspond to different steps of the computation, for every input there is a value of i such that the i -CS is short, namely has length at most $t(|x|)/n$. If $t(n)$ is much less than n^2 , the length of this CS is much less than n , from which we can conclude that the number of CSs is much less than the number of inputs, and so we can find two inputs with the same CS.

Proof. Let n be divisible by four, without loss of generality, and consider palindromes of the form

$$p(x) := x0^{n/2}x^R$$

where $x \in [2]^{n/4}$ and x^R is the reverse of x .

Assume there are $x \neq y$ in $[2]^{n/4}$ and i in the middle part, i.e., $n/4 \leq i \leq 3n/4 - 1$, so that the i -CS of M on $p(x)$ and $p(y)$ is the same. Then we can define $z := x0^{n/2}y^R$ which is not a palindrome but is still accepted by M , concluding the proof.

There remains to prove that the assumption of Theorem 3.1 implies the assumption in the previous paragraph. Suppose M runs in time t . Since crossing sequences at different boundaries correspond to different steps of the computation, for every $x \in [2]^{n/4}$ there is a value of i in the middle part such that the i -CS of M on $p(x)$ has length $\leq ct/n$. This implies that there is an i in the middle s.t. there are $\geq c2^{n/4}/n$ inputs x for which the i -CS of M on x has length $\leq ct/n$.

For fixed i , the number of i -CS of length $\leq \ell$ is $\leq (s+1)^\ell$.

Hence there are $x \neq y$ for which $p(x)$ and $p(y)$ have the same i -CS whenever $c2^{n/4}/n \geq (s+1)^{ct/n}$. Taking logs one gets $ct \log(s)/n \leq cn$. **QED**

Exercise 3.1. For every s and n describe an s -state TM deciding palindromes in time $cn^2/\log s$ (matching Theorem 3.1).

Exercise 3.2. Let $L := \{xx : x \in [2]^*\}$. Show $L \in \text{TM-Time}(cn^2)$, and prove this is tight up to constants.

One may be tempted to think that it is not hard to prove stronger bounds for similar functions. In fact as mentioned above this has resisted all attempts!

3.2 Counting: impossibility results for non-explicit functions

Proving the *existence* of hard functions is simple: Just count. If there are more functions than efficient machines, some function is not efficiently computable. This is applicable to any model; next we state it for TMs for concreteness. Later we will state it for circuits.

Theorem 3.2. There exists a function $f : [2]^n \rightarrow [2]$ that cannot be computed by a TM with s states unless $cs \log s \geq 2^n$, regardless of time.

Proof. The number of TMs with s states is $\leq s^{cs}$, and each TM computes at most one function (it may compute none, if it does not stop). The number of functions on n bits is 2^{2^n} . Hence if $2^n > cs \log s$ some function cannot be computed. **QED**

Note this bound is not far from that in Exercise 2.3.

It is instructive to present this basic result as an application of the probabilistic method:

Proof. Let us pick f uniformly at random. We want to show that

$$\mathbb{P}_f[\exists \text{ an } s\text{-state TM } M \text{ such that } M(x) = f(x) \text{ for every } x \in [2]^n] < 1.$$

Indeed, if the probability is less than 1 than some function exists that cannot be computed. By a union bound we can say that this probability is

$$\leq \sum_M \mathbb{P}_f[M(x) = f(x) \text{ for every } x \in [2]^n],$$

where the sum is over all s -state machines. Each probability in the sum is $(1/2)^{2^n}$, since M is fixed. The number of s -state machines is $\leq s^{cs}$. So the sum is $\leq s^{cs}(1/2)^{2^n}$, and we can conclude as before taking logs. **QED**

3.3 Diagonalization and time hierarchy

Can you compute more if you have more time? For example, can you write a program that runs in time n^2 and computes something that cannot be computed in time $n^{1.5}$? The answer is yes for trivial reasons if we allow for non-boolean functions.

Exercise 3.3. Give a function $f : [2]^* \rightarrow [2]^*$ in $\text{Time}(n^2) \setminus \text{Time}(n^{1.5})$.

The answer is more interesting if the functions are boolean. Such results are known as *time hierarchies*, and a generic technique for proving them is *diagonalization*, applicable to any model.

We first illustrate the result in the simpler case of partial functions, which contains the main ideas. Later we discuss total functions.

Theorem 3.3. There is a partial function in $\text{TM-Time}(t(n))$ such that any TM M computing it runs in time $\geq c_M t(n)$, for any $t(n) = \omega(1)$.

In other words, $\text{Time}(t(n)) \not\supseteq \text{Time}(o(t(n)))$.

Proof. Consider the TM H that on input $x = (M, 1^{n-|M|})$ of length n runs M on x until it stops and then complements the answer. (We can use a simple encoding of these pairs, for example every even-position bit of the description of M is a 0.)

Now define X to be the subset of pairs s.t. M runs in time $\leq t(n)/|M|^c$ on inputs of length n , and $|M|^c \leq t(n)/2$.

On these inputs, H runs in time $|M|^c + |M|^c \cdot t(n)/|M|^c \leq t(n)$, as desired. To accomplish this, H can begin by making a copy of M in time $|M|^c \leq t(n)/2$. Then every step of the computation of M can be simulated by H with $|M|^c$ steps, always keeping the description of M to the left of the head.

Now suppose N computes the same function as H in time $t(n)/|N|^c$. Note that $x := (N, 1^{n-|N|})$ falls in the domain X of the function, for n sufficiently large, using that $t(n) = \omega(1)$. Now consider running N on x . We have $N(x) = H(x)$ by supposition, but $H(x)$ is the complement of $N(x)$, contradiction. **QED**

This proof is somewhat unsatisfactory; in particular we have no control on the running time of H on inputs not in X . It is desirable to have a version of this fundamental result

for total functions. Such a version is stated next. It requires additional assumptions on t and a larger gap between the running times. Perhaps surprisingly, as we shall discuss, both requirements are essential.

Theorem 3.4. Let $t(n) \geq n$ be a function. Suppose that $f(x) := t(|x|)$ is in $\text{TM-Time}(t(n)/\log^c n)$.

There is a total function in $\text{TM-Time}(ct(n)\log t(n))$ that cannot be computed by any TM M in time $c_M t(n)$.

The assumption about t is satisfied by all standard functions, including all those in this book. (For example, take $t(n) := n^2$. The corresponding f is then $|x|^2$. To compute f on input x of n bits we can first compute $|x| = n$ in time $cn \log n$ (Exercise 2.2). This is a number of $b := \log n$ bits. We can then square this number in time b^c . Note that the time to compute $f(x)$ is dominated by the $cn \log n$ term coming from computing $|x|$, which does not depend on t and is much less than the $n^2/\log^c n$ in the assumption.) The assumption cannot be removed altogether because there exist pathological functions t for which the result is false.

The proof is similar to that of Theorem 3.3. However, to make the function total we need to deal with arbitrary machines, which may not run in the desired time or even stop at all. The solution is to clock H in a manner similar to the proof of the universal machine, Lemma 2.1.

Also, we define a slightly different language to give a stronger result – a unary language – and to avoid some minor technical details (the possibility that the computation of f erases the input).

We define a TM H that on input 1^n obtains a description of a TM M , computes $t(n)$, and then simulates M on input 1^n for $t(n)$ steps in a way similar to Lemma 2.1, and if M stops then H outputs the complement of the output of M ; and if M does not stop then H stops and outputs anything. Now H computes a function in time about $t(n)$. We argue that this function cannot be computed in much less time as follows. Suppose some TM M computes the function in time somewhat less than $t(n)$. Then we can pick an 1^n for which H obtains the description of this M , and the simulation always stops. Hence, for that 1^n we would obtain $M(1^n) = H(1^n) = 1 - M(1^n)$, which is a contradiction.

However, there are interesting differences with the simulation in Lemma 2.1. In that lemma the universal machine U was clocking the steps of the machine M being simulated. Now instead we need to clock the steps of U itself, even while U is parsing the description of M to compute its transition function. This is necessary to guarantee that H does not waste time on big TM descriptions.

Whereas in Lemma 2.1 the tape was arranged as

$$(x, M, \dot{i}, t', y),$$

it will now be arranged as

$$(x, M', \dot{i}, t', M'', y)$$

which is parsed as follows. The description of M is $M'M''$, M is in state \dot{i} , the tape of M contains xy and the head is on the left-most symbol of y . The integer t' is the counter decreased at every step

Proof. Define TM H that on input 1^n :

1. Compute $(n, t(n), 1^n)$.
2. Compute $(M_n, t(n), 1^n)$. Here M_n is obtained from n by removing all left-most zeroes until the first 1. I.e., if $n = 0^j 1x$ then $M_n = x$. This is similar to the fact that a program does not change if you add, say, empty lines at the bottom.
3. Simulate M_n on 1^n , reducing the counter $t(n)$ at every step, including those parsing M_n , as explained before.
4. If M_n stops before the counter reaches 0, output the complement of its output. If the counter reaches 0 stop and output anything.

Running time of H .

1. Computing n is similar to Exercise 2.2. By assumption $t(n)$ is computable in time $t(n)/\log^c n$. Our definition of computation allows for erasing the input, but we can keep n around spending at most another $\log^c n$ factor. Thus we can compute $(n, t(n))$ in time $t(n)$. Finally, in case it was erased, we can re-compute 1^n in time $cn \log n$ by keeping a counter (initialized to a copy of n).
2. This takes time $c(n + t(n))$: simply scan the input and remove zeroes.
3. Decreasing the counter takes $c|t(n)|$ steps. Hence this simulation will take $ct(n) \log t(n)$ time.

Overall the running time is $ct(n) \log t(n)$.

Proof that the function computed by H requires much time. Suppose some TM M computes the same function as H . Consider inputs 1^n where $n = 0^j 1M$. Parsing the description of M to compute its transition function takes time c_M , a value that depends on M only and not on j . Hence H will simulate $\lfloor t(n)/c_M \rfloor$ steps of M . If M stops within that time (which requires $t(n)$ to be larger than c_M , and so n and j sufficiently large compared to M) then the simulation terminates and we reach a contradiction as explained before. **QED**

The extra $\log t(n)$ factor cannot be reduced because of the surprising result presented in Theorem 3.5 showing that, on TMs, time $o(n \log n)$ equals time n for computing total functions.

However, tighter time hierarchies hold for more powerful models, like RAMs. Also, a time hierarchy for total functions for BPTIME is... an open problem! But a hierarchy is known for partial functions.

Exercise 3.4. (1) State and prove a tighter time hierarchy for Time (which recall corresponds to RAMs) for total functions. You don't need to address simulation details, but you need to explain why a sharper separation is possible.

(2) Explain the difficulty in extending (1) to BPTIME.

(3) State and prove a time hierarchy for BPTIME for partial functions.

3.3.1 TM-Time($o(n \log n)$) = TM-Time(n)

In this subsection we prove the result in the title, which we also mentioned earlier. First let us state the result formally.

Theorem 3.5. Let $f : [2]^* \rightarrow [2]$ be in TM-Time($t(n)$) for a $t(n) = o(n \log n)$. Then $f \in \text{TM-Time}(n)$.

Note that time n is barely enough to scan the input. Indeed, the corresponding machines in Theorem 3.5 will only move the head in one direction.

The rest of this section is devoted to proving the above theorem. Let M be a machine for f witnessing the assumption of the theorem. We can assume that M stops on every input (even though our definition of time only applies to large enough inputs), possibly by adding $\leq n$ to the time, which does not change the assumption on $t(n)$. The theorem now follows from the combination of the next two lemmas.

Lemma 3.1. Let M be a TM running in time $t(n) \leq o(n \log n)$. Then on every input $x \in [2]^*$ every i -CS with $i \leq |x|$ has length $\leq c_M$.

Proof. Assume towards a contradiction that for every $b \in \mathbb{N}$ there are inputs which have crossing sequences of length $\geq b$. Specifically let $x(b)$ be a shortest input of length $n(b) := |x(b)|$ such that there exists $j \in \{0, 1, \dots, n(b)\}$ for which the j -CS in the computation of M on $x(b)$ has length $\geq b$.

Exercise 3.5. Prove $n(b) \rightarrow \infty$ for $b \rightarrow \infty$.

There are $n(b)+1 \geq n(b)$ tape boundaries within or bordering $x(b)$. If we pick a boundary uniformly at random, the average length of a CS on $x(b)$ is $\leq t(n(b))/n(b)$. Hence there are $\geq n(b)/2$ choices for i s.t. the i -CS on $x(b)$ has length $\leq 2t(n(b))/n(b)$.

The number of such crossing sequences is

$$\leq (s+1)^{2t(n(b))/n(b)} = (s+1)^{o(n(b) \log(n(b)))/n(b)} = n(b)^{o(\log s)}.$$

Hence, the same crossing sequence occurs at $\geq (n(b)/2)/n(b)^{o(\log s)} \geq 4$ positions i , using that $n(b)$ is large enough.

Of these four, one could be the CS of length $\geq b$ from the definition of $x(b)$. Of the other three, two are on the same side of j . We can remove the corresponding interval of the input without removing the CS of length $\geq b$. Hence we obtained a shorter input with a CS of length $\geq b$, contradicting our definition of $x(b)$ and so our initial assumption. **QED**

Lemma 3.2. Suppose $f : [2]^* \rightarrow [2]$ is computable by a TM such that on every input x , every i -CS with $i \leq |x|$ has length $\leq b$. Then f is computable in time n by a TM with c_b states that only moves the head in one direction.

Exercise 3.6. Prove this.

3.4 Circuits

The situation for circuits is similar to that of 2-TMs, and by Theorem 2.6 we know that proving $\omega(n \log n)$ bounds on circuits is harder than for 2-TMs. Even a bound of cn is unknown. The following is a circuit version of Theorem 3.2. Again, the bound is close to what we saw in Theorem 2.4.

Theorem 3.6. There are functions $f : [2]^n \rightarrow [2]$ that require circuits of size $\geq (1 - o(1))2^n/n$, for every n .

One can prove a hierarchy for circuit size, by combining Theorem 3.6 and Theorem 2.4. As stated, the results give that circuits of size cs compute more functions than those of size s . In fact, the “ $o(1)$ ” in the theorems is small, so one can prove a sharper result. But a stronger and more enjoyable argument exists.

Theorem 3.7. [Hierarchy for circuit size] For every n and $s \leq c2^n/n$ there is a function $f : [2]^n \rightarrow [2]$ that can be computed by circuits of size $s + cn$ but not by circuits of size s .

Proof. Consider a path from the all-zero function to a function which requires circuits of size $\geq s$, guaranteed to exist by Theorem 3.6, changing the output of the function on one input at each step of the path. Let h be the first function that requires size $> s$, and let h' be the function right before that in the path. Note that h' has circuits of size $\leq s$, and h can be computed from h' by changing the value on a single input. The latter can be implemented by circuits of size cn . **QED**

Exercise 3.7. Prove a stronger hierarchy result for alternating circuits, where the “ cn ” in Theorem 3.7 is replaced with “ c .”

In fact, this improvement is possible even for (non alternating) circuits, see Problem 3.2.

3.4.1 The circuit won't fit in the universe: Non-asymptotic, cosmological results

Most of the results in this book are *asymptotic*, i.e., we do not explicitly work out the constants because they become irrelevant for larger and larger input lengths. As stated, these results don't say anything for inputs of a fixed length. For example, any function on 10^{100} bits is in $\text{Time}(c)$.

However, it is important to note that all the proofs are *constructive* and one can work out the constants, and produce non-asymptotic results. We state next one representative example when this was done. It is about a problem in logic, an area which often produces very hard problems.

On an alphabet of size 63, the language used to write formulas includes first-order variables that range over \mathbb{N} , second-order variables that range over finite subsets of \mathbb{N} , the predicates “ $y = x + 1$ ” and “ $x \in S$ ” where x and y denote first-order variables and S denotes

a set variable, and standard quantifiers, connectives, constants, binary relation symbols on integers, and set equality. For example one can write things like “every finite set has a maximum:” $\forall S \exists x \in S \forall y \in S, y \leq x$.

Theorem 3.8. [181] To decide the truth of logical formulas of length at most 610 in this language requires a circuit containing at least 10^{125} gates. So even if each gate were the size of a proton, the circuit would not fit in the known universe.

Their result applies even to randomized circuits with error $1/3$, if 610 is replaced with 614. (We can define randomized circuits analogously to BPTIME.)

3.5 Problems

Problem 3.1. Hierarchy Theorem 3.4 only gives a function f that cannot be computed fast on *all* large enough input lengths: it is consistent with the theorem that f can be computed fast on infinitely many input lengths.

Give a function $f : [2]^* \rightarrow [2]^*$ mapping x to $[2]^{\log \log \log |x|}$ that is computable in time n^c but such that for any RAM M running in time n^2 the following holds. For every $n \geq c_M$ and every $x \in [2]^n$ one has $M(x) \neq f(x)$.

Hint: Note the range of f . Can this result hold as stated with range $[2]$?

Problem 3.2. Replace “ cn ” in Theorem 3.7 with “ c .”

Problem 3.3. Prove that $\{0^i 1^i : i \geq 0\} \in \text{TM-Time}(cn \log n) \setminus \text{TM-Time}(t(n))$, for any $t(n) = o(n \log n)$.

For the negative result, don’t use pumping lemmas or other characterization results not covered in this book.

Problem 3.4. The following argument contradicts Theorem 3.4; what is wrong with it?

“By Theorem 3.5, $\text{TM-Time}(n \log^{0.9} n) = \text{TM-Time}(n)$. By padding (Theorem 3.5), $\text{TM-Time}(n \log^{1.1} n) = \text{TM-Time}(n \log^{0.9} n)$. Hence $\text{TM-Time}(n \log^{1.1} n) = \text{TM-Time}(n)$.”

3.6 Notes

Concluding, I view the mystery of the difficulty of proving (even the slightest non-trivial) computational difficulty of natural problems to be one of the greatest mysteries of contemporary mathematics. [219]

Crossing sequences and the tight quadratic bound for palindromes are from [89].

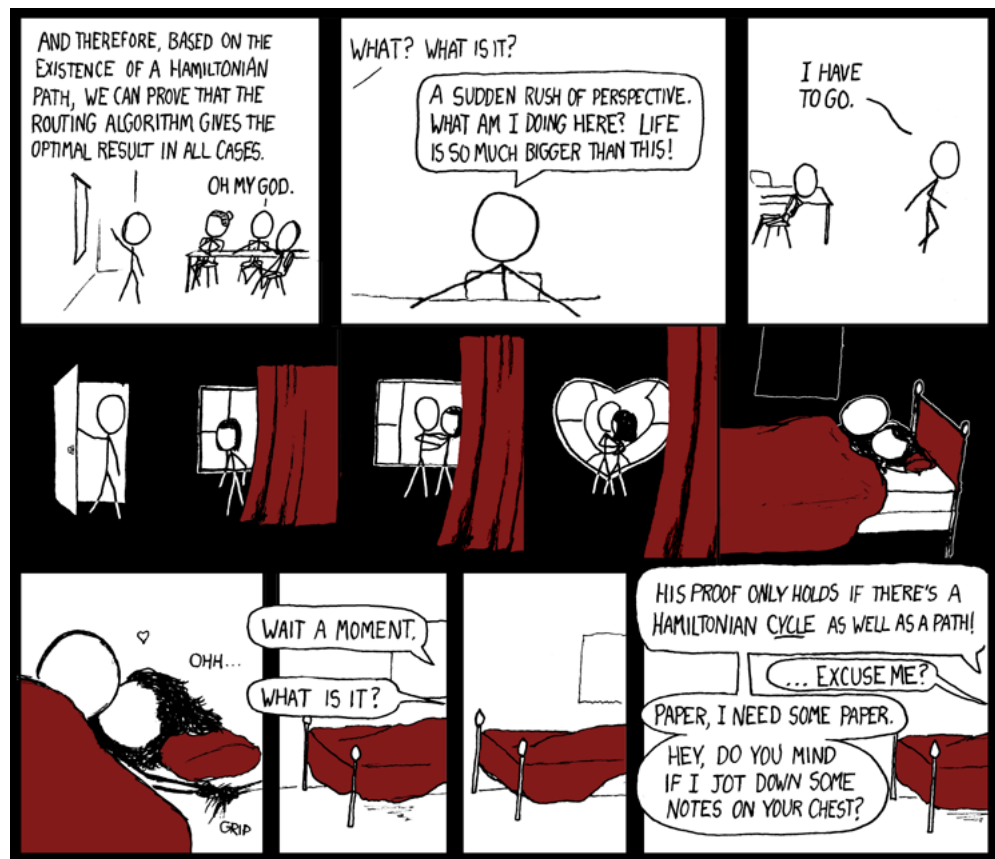
The time hierarchy originates in [91] and was later optimized [92].

The existence of hard functions via counting arguments, Theorem 3.6, goes back to [172], Theorem 7, “*Are most functions simple or complex?*”

Theorem 3.5 follows by combining results in [90, 114].

Chapter 4

Reductions



<https://xkcd.com/230/>

One can relate the complexity of functions via *reductions*. This concept is so ingrained in common reasoning that giving it a name may feel, at times, strange. For in some sense pretty much everything proceeds by reductions. In any algorithms textbook, the majority of algorithms can be cast as reductions to algorithms presented earlier in the book, and so on. And it is worthwhile to emphasize now that, as we shall see below, reductions, even in the

context of computing, have been used for millennia. For about a century reductions have been used in the context of undecidability in a modern way, starting with the incompleteness theorem in logic, whose proof reduces questions in logic to questions in arithmetic.

Perhaps one reason for the more recent interest in complexity reductions is that we can use them to relate problems that are tantalizingly close to problems that today we solve routinely on somewhat large scale inputs with computers, and that therefore appear to be just out of reach. By contrast, reductions in the context of undecidability tend to apply to problems that are completely out of reach, and in this sense remote from our immediate worries.

4.1 Types of reductions

Informally, a reduction from a function f to a function g is a way to compute f given that we can compute g . One can define reductions in different ways, depending on the overhead required to compute f given that we can compute g . The most general type of reduction is simply an *implication*.

General form of reduction from f to g :
If g can be computed with resources X then f can be computed with resources Y .

A common setting is when $X = Y$. In this case the reduction allows us to stay within the same complexity class.

Definition 4.1. We say that f reduces to g in X (or under X reductions) if

$$g \in X \Rightarrow f \in X.$$

A further special and noteworthy case is when $X = P$, or $X = BPP$; in these cases the reduction can be interpreted as saying that if g is easy to compute then f is too. But in general X may not be equal to Y . We will see examples of such implications for various X and Y .

It is sometimes useful to be more specific about how the implication is proved. For example, this is useful when inferring various properties of f from properties of g , something which can be obscured by a stark implication. The following definition gives a specific way in which the implication can be proved.

Definition 4.2. We say that f map reduces to g in X (or via a map in X) if there is $M \in X$ such that $f(x) = g(M(x))$ for every x .

Exercise 4.1. Suppose that f map reduces to g in X .

- (1) Suppose $X = P$. Show f reduces to g in X .
- (2) Suppose $X = \bigcup_d \text{Time}(d \cdot n^2)$. Can you still show that f reduces to g in X ?

In general, the harder the problems we are reducing the more constrained the class X can be. In several interesting cases, extremely constrained classes X suffice for the reductions, such as functions with constant locality, or even locality one, cf. Problem 4.5. We define this notion of locality next, which is important to keep in mind and will play a role in later chapters as well.

Definition 4.3. A function $f : [2]^n \rightarrow [2]^m$ is d -local if each output bit depends on $\leq d$ input bits.

But for many other problems, the reductions we shall see are not even mapping reductions. In fact, our first example is not a mapping reduction.

4.2 Reductions

4.2.1 Multiplication

Summing two n -bit integers is in $\text{CktGates}(cn)$ (Exercise 2.8). But the smallest circuit known for multiplication has $\geq cn \log n$ gates. (The same situation holds for MTMs; over RAMs and related models multiplication can be done in time cn .) It is a long-standing question whether we can multiply two n -bit integers with a linear-size circuit.

What about squaring integers? Is that harder or easier than multiplication? Obviously, if we can multiply two numbers we can also square a number: simply multiply it by itself. This is a trivial example of a reduction. What about the other way around? We can use a reduction established millennia ago by the Babylonians. They employed the equation

$$a \cdot b = \frac{(a + b)^2 - (a - b)^2}{4} \quad (4.1)$$

to reduce multiplication to squaring, plus some easy operations like addition and division by four. In our terminology we have the following.

Definition 4.4. Multiplication is the problem of computing the product of two n -bit integers. Squaring is the problem of computing the square of an n -bit integer.

Theorem 4.1. If Squaring has linear-size circuits then Multiplication has linear-size circuits.

Proof. Suppose C computes Squaring. Then we can multiply using equation (4.1). Specifically, given a and b we use Exercise 2.8 to compute $a + b$ and $a - b$. (We haven't seen subtraction or negative integers, but it's similar to addition.) Then we run C on both of them. Finally, we again use Exercise 2.8 for computing their difference. It remains to divide by four. In binary, this is accomplished by ignoring the last two bits – which costs nothing on a circuit. **QED**

4.2.2 3Sum

Definition 4.5. The 3Sum problem: Given a list of integers, are there three integers that sum to 0?

It is easy to solve 3Sum in time $cn^2 \log n$ on a RAM. (We can first sort the integers then for each pair (a, b) we can do a binary search to check if $-(a + b)$ is also present.) The time can be improved $n^2 / \log^c n$.

3Sum is believed to require quadratic time.

Definition 4.6. $\text{SubquadraticTime} := \bigcup_{\epsilon > 0} \text{Time}(n^{2-\epsilon})$.

Conjecture 4.1. $3\text{Sum} \notin \text{SubquadraticTime}$.

One can reduce 3Sum to a number of other interesting problem to infer that, under Conjecture 4.1, those problems require quadratic time too.

Definition 4.7. The Collinearity problem: Given a list of points in the plane, are there three points on a line?

Theorem 4.2. $\text{Collinearity} \in \text{SubquadraticTime} \Rightarrow 3\text{Sum} \in \text{SubquadraticTime}$ (i.e., Conjecture 4.1 is false).

Proof. We map instance a_1, a_2, \dots, a_n of 3Sum to the points

$$(a_1, a_1^3), (a_2, a_2^3), \dots, (a_n, a_n^3),$$

and solve Collinearity on those points.

To verify correctness, notice that points (x, x^3) , (y, y^3) , and (z, z^3) are on a line iff

$$\frac{y^3 - x^3}{y - x} = \frac{z^3 - x^3}{z - x}.$$

Because $y^3 - x^3 = (y - x)(y^2 + yx + x^2)$, this condition is equivalent to

$$y^2 + yx + x^2 = z^2 + zx + x^2 \Leftrightarrow (x + (y + z))(y - z) = 0.$$

Assuming $y \neq z$, i.e., that the 3Sum instance consists of distinct numbers, this is equivalent to $x + y + z = 0$, as desired. (The case where there can be duplicates is left as an exercise.)

Note that the Collinearity instance has length linear in the 3Sum instance, and the result follows. **QED**

Exercise 4.2. The Tripartite-3Sum problem: Given lists A_1 , A_2 , and A_3 of numbers, are there $a_i \in A_i$ s.t. $a_1 + a_2 + a_3 = 0$?

Prove that Tripartite-3Sum is in subquadratic time iff 3Sum is.

We now give a reduction in the other direction: We reduce a problem to 3Sum.

Definition 4.8. The 3Cycle-Detection problem: Given the adjacency list of a directed graph, is there a cycle of length 3?

This problem can be solved in time $n^{2\omega/(\omega+1)+o(1)}$ where $\omega < 2.373$ is the exponent of matrix multiplication. If $\omega = 2$ then the bound is $n^{1.33+o(1)}$. It is not known if any subquadratic algorithm for 3Sum would improve these bounds. However, we can show that an improvement follows if $3\text{Sum} \in \text{Time}(n^{1+\epsilon})$ for a small enough ϵ .

Theorem 4.3. $3\text{Sum} \in \text{Time}(t(n)) \Rightarrow 3\text{Cycle-Detection} \in \text{BPTime}(ct(n))$, for any $t(n) \geq n$.

The reduction can be derandomized (that is, one can replace BPTime with Time in the conclusion) but the randomized case contains the main ideas.

Proof. We assign random numbers r_x with $4 \log n$ bits to each node x in the graph. The 3Sum instance consists of the integers $r_x - r_y$ for every edge $x \rightarrow y$ in the graph.

To verify correctness, suppose that there is a cycle

$$x \rightarrow y \rightarrow z \rightarrow x$$

in the graph. Then we have $r_x - r_y + r_y - r_z + r_z - r_x = 0$, for any random choices.

Conversely, suppose there is no cycle, and consider any three numbers $r_{x1} - r_{y1}, r_{x2} - r_{y2}, r_{x3} - r_{y3}$ from the reduction and its corresponding edges. Some node x_i has unequal in-degree and out-degree in those edges. This means that when summing the three numbers, the random variable r_{x_i} will not cancel out. When selecting uniform values for that variable, the probability of getting 0 is at most $1/n^4$.

By a union bound, the probability there are three numbers that sum to zero is $\leq n^3/n^4 < 1/3$. **QED**

Exercise 4.3. Prove analogous results for:

3Sum vs. 3 Cycles on undirected graphs.

4Sum vs. 4 Cycles on undirected graphs. Hint: This might not be as easy as the first part.

To be clear, in the input to the undirected graph problems we do not allow repeated edges among nodes, and a cycle cannot use an edge more than once.

Many other clusters of problems exist, for example based on matrix multiplication or all-pairs shortest path.

4.3 Reductions from 3Sat

In this section we begin to explore an important cluster of problems not known to be in BPP. What's special about these problems is that in Chapter 5 we will show that we can reduce *arbitrary computation* to them, while this is unknown for the problems in the previous section.

Perhaps the most basic problem in the cluster is the following.

Definition 4.9. A 3CNF is a CNF where every clause has at most three literals. The 3Sat problem: Given a 3CNF ϕ , is there an assignment x s.t. $\phi(x) = 1$?

Conjecture 4.2. 3Sat \notin P.

Stronger conjectures have been made.

Conjecture 4.3. [Exponential time hypothesis (ETH)] There is $\epsilon > 0$ such that there is no algorithm that on input a 3CNF ϕ with v variables and cv^3 clauses decides if ϕ is satisfiable in time $2^{(\epsilon+o(1))v}$.

Conjecture 4.4. [Strong exponential-time hypothesis (SETH)] For every $\epsilon > 0$ there is k such that there is no algorithm that on input a k CNF ϕ with v variables and cv^k clauses decides if ϕ is satisfiable in time $2^{(1-\epsilon+o(1))v}$.

It is known that SETH \Rightarrow ETH, but the proof is not immediate.


We now give reductions from 3Sat to several other problems. The reductions are in fact mapping reductions. Moreover, the reduction map can be extremely restricted, see Problem 4.5. In this sense, therefore, these reductions can be viewed as direct translations of the problems, and maybe we shouldn't really be thinking of the problems as different, even if they at first sight refer to different objects (formulas, graphs, numbers, etc.).

For videos covering these reductions you can watch videos 29, 30, 31, and 32 covering reductions: 3SAT to CLIQUE, CLIQUE to VERTEX-COVER, 3SAT to SUBSET-SUM, 3SAT to 3COLOR from <https://www.ccs.neu.edu/home/viola/classes/algm-generic.html> Note: The videos use the terminology "polynomial time" instead of "power time" here.

4.3.1 3Sat to Clique

Definition 4.10. The Clique problem, given a graph G and an integer t , are there t nodes in G that are all connected? The latter is called a *clique* of size t .

Example 4.1. The following graph has a clique of size 3 but not of size 4:



Theorem 4.4. Clique \in P \Rightarrow 3Sat \in P.

Proof. Given a 3CNF φ with k clauses, we construct a graph G with $3k$ nodes where we have a node for each literal occurrence. We then connect all except

- (A) Nodes in same clause, and
- (B) Contradictory nodes, such as x and $\neg x$.

The construction is in P.

We claim that φ is satisfiable iff G has a clique of size k .


Only if: Given a satisfying assignment, collect exactly one node which is satisfied in each clause. This makes $t = k$ nodes. For any pair of such nodes, (A) does not hold by construction, and (B) because they correspond to an assignment.

If: Given a clique of size t , pick any assignment that makes the corresponding literals true. This is a valid definition by (B). Also, because of (A), there is at least one true literal in each clause. **QED**

Example 4.2. Consider


$$\varphi = (x \vee y \vee z) \wedge (\neg x \vee \neg y \vee z) \wedge (x \vee y \vee \neg z).$$

The corresponding graph G is:




We seek cliques of size $t = k = 3$, a.k.a. triangles.

A satisfying assignment to φ is $x = 0; y = 1; z = 0$. The corresponding clique is shown next:



Another satisfying assignment is $x = 1; y = 0; z = 1$. The corresponding clique is shown next:



4.3.2 Clique to cover-by-vertexes

TBD

4.3.3 3Sat to Subset-Sum

Definition 4.11. The Subset-sum problem: Given n integers a_i and a target t , is there a subset of the a_i that sums to t ?

Example 4.3. There is a subset of 5, 2, 14, 3, 9 summing to $t := 25$ ($2 + 14 + 9 = 25$). But there is no subset of 1, 3, 4, 9 summing to $t := 15$.

Subset-sum is also a very interesting problems. If the numbers are small it can be solved in power time via dynamic programming. Hence the next reduction capitalizes on the magnitude of the numbers.

Theorem 4.5. Subset-sum $\in P \Rightarrow 3\text{Sat} \in P$.

Proof. On input φ with v variables and k clauses we produce a list of numbers with $v + k$ digits. The most significant v correspond to variables; the other k to clauses. For each variable x include number a_x^T which has 1 in the digit corresponding to x , and a 1 in every digit of a clause where x appears without negation. Similarly, include number a_x^F which also has a 1 in the digit corresponding to x , and now a 1 in every digit of a clause where x appears negated.

Also, for each clause C , include twice the number a_C which has a 1 in the digit corresponding to C , 0 in others.

Set t to be 1 in first v digits, and 3 in rest k digits.

This construction is power time.

Now suppose φ has satisfying assignment. Pick a_x^T if x is true, a_x^F if x is false. The sum of these numbers yield 1 in first v digits by construction. It also yields 1, 2, or 3 in each of the last k digits because each clause has a true literal. By picking appropriate subset of the numbers a_C we can reach t .

Conversely, given a subset, note that there is no carry in sum, because there are only 3 literals per clause. So digits behave “independently.” For each pair a_x^T, a_x^F exactly one is included, otherwise would not get 1 in that digit. Define x true if a_x^T included, false otherwise. For any clause C , the a_C contribute ≤ 2 in that digit. So each clause must have a true literal otherwise sum would not get to 3 in that digit. **QED**

Example 4.4. Let $\varphi := (x \vee y \vee z) \wedge (\neg x \vee \neg y \vee z) \wedge (x \vee y \vee \neg z)$. The subset-sum instance is:

	var x	var y	var z	clause 1	clause 2	clause 3
$a_x^T =$	1	0	0	1	0	1
$a_x^F =$	1	0	0	0	1	0
$a_y^T =$	0	1	0	1	0	1
$a_y^F =$	0	1	0	0	1	0
$a_z^T =$	0	0	1	1	1	0
$a_z^F =$	0	0	1	0	0	1
$a_{c1} =$	0	0	0	1	0	0
$a_{c2} =$	0	0	0	0	1	0
$a_{c3} =$	0	0	0	0	0	1
$t =$	1	1	1	3	3	3

A satisfying assignment is $x = 0, y = 1, z = 0$. The corresponding subset is:

	var x	var y	var z	clause 1	clause 2	clause 3
$a_x^T =$	1	0	0	1	0	1
$a_x^F =$	1	0	0	0	1	0
$a_y^T =$	0	1	0	1	0	1
$a_y^F =$	0	1	0	0	1	0
$a_z^T =$	0	0	1	1	1	0
$a_z^F =$	0	0	1	0	0	1
(2x) $a_{c1} =$	0	0	0	1	0	0
(2x) $a_{c2} =$	0	0	0	0	1	0
(2x) $a_{c3} =$	0	0	0	0	0	1
$t =$	1	1	1	3	3	3

(choose twice)
(choose twice)

Another satisfying assignment is $x = y = z = 1$ with corresponding subset


	var x	var y	var z	clause 1	clause 2	clause 3
$a_x^T =$	1	0	0	1	0	1
$a_x^F =$	1	0	0	0	1	0
$a_y^T =$	0	1	0	1	0	1
$a_y^F =$	0	1	0	0	1	0
$a_z^T =$	0	0	1	1	1	0
$a_z^F =$	0	0	1	0	0	1
(2x) $a_{c1} =$	0	0	0	1	0	0
(2x) $a_{c2} =$	0	0	0	0	1	0
(2x) $a_{c3} =$	0	0	0	0	0	1
$t =$	1	1	1	3	3	3


(choose twice)

4.3.4 3Sat to 3Color

Definition 4.12. A 3-coloring of a graph is a coloring of each node, using at most 3 colors, such that no adjacent nodes have the same color. The 3Color problem: Given a graph G , does it have a 3 coloring?

Example 4.5. The following graphs have a 3-coloring, shown:





An example of a graph that cannot be 3-colored is a clique of size 4.

Theorem 4.6. $3\text{Color} \in \text{P} \Rightarrow 3\text{Sat} \in \text{P}$.


Proof. Given a 3CNF ϕ , we construct a graph G as follows.

Add 3 special nodes called the "palette" in a clique:


T = True

F = False


B = Base




For each variable add 2 literal nodes with an edge between them




For each clause add the following gadget with 6 nodes



Connect each literal node to node B in the palette



For each clause (l_1, l_2, l_3) connect the clause gadget to the palette and to the nodes l_i as follows:




The construction of G is in P. We now prove that φ is satisfiable iff G is 3 colorable. We begin with some preliminary remarks. In the palette, T's color represents TRUE, and F's color represents FALSE. Note in a 3-coloring, all variable nodes must be colored T or F because they are connected to B. Also, x and $\neg x$ must have different colors because they are connected. So we can “translate” a 3-coloring of G into a true/false assignment to variables of φ .


The important claim is that a clause gadget can be 3-colored iff any of the literals connected to it is colored True. This holds because each of the two triangles in a the clause gadget is computing “Or.” In a triangle, the top node is colored according to the Or of the two literals connected to the bottom two nodes in the triangle. For example, if the literals are both F, then the bottom nodes in the triangle must be colored T and B, and so the top is F.

The result follows. Given a satisfying assignment, we can pick the corresponding coloring of the literal nodes and extend it to a 3 coloring of the entire graph. Vice versa, given a 3 coloring of the graph we can infer an assignment to the variables and note that each clause has a true literal since each clause gadget is 3 colored. **QED**

Example 4.6. Let $\varphi := (x \vee y \vee z) \wedge (\neg x \vee \neg y \vee z)$. Then G is



A satisfying assignment is $x = y = 0$ and $z = 1$. The corresponding coloring is



Exercise 4.4. The problem System is defined as follows. A *linear inequality* is an inequality involving sums of variables and constants, such as $x + y \geq z$, $x \leq -17$, and so on. A system of linear inequalities has an *integer* solution if it is possible to substitute integer values for the variables so that every inequality in the system becomes true. The language System consists of systems of linear inequalities that have an integer solution. For example,

$$(x + y \geq z, x \leq 5, y \leq 1, z \geq 5) \in \text{System}$$

$$(x + y \geq 2z, x \leq 5, y \leq 1, z \geq 5) \notin \text{System}$$

Reduce 3Sat to System in P.

Exercise 4.5. For an integer k , k -Color is the problem of deciding if the nodes of a given undirected graph G can be colored using k colors in such a way that no two adjacent vertices have the same color.

Reduce 3-Color to 4-Color in P.

Reductions in the opposite directions are possible, and so in fact the problems in this section are *power-time equivalent* in the sense that any of the problems is in P iff all the others are. We will see a generic reduction in the next chapter. For now, we illustrate this equivalence in a particular case.

Exercise 4.6. Reduce 3Color to 3Sat in P, following these steps:

1. Given a graph G , introduce variables $x_{i,d}$ representing that node i has color d , where d ranges in the set of colors $C = \{g, r, b\}$. Describe a set of clauses that is satisfiable if and only if for every i there is exactly one $d \in C$ such that $x_{i,d}$ is true.
2. Introduce clauses representing that adjacent nodes do not have the same color.
3. Briefly conclude the proof.

Thus, we are identifying a cluster of problems which are all power-time equivalent.

4.4 Power hardness from SETH

In this section we show that a conjecture similar to Conjecture 4.1 can be proved assuming SETH. This is an interesting example of how we can connect different parameter regimes, since SETH is stated in terms of exponential running times. In general, “scaling” parameters is a powerful technique in the complexity toolkit.

Definition 4.13. The Or-Vector problem: Given two sets A and B of strings of the same length, determine if there is $a \in A$ and $b \in B$ such that the bit-wise Or $a \vee b$ equals the all-one vector.

The Or-Vector problem is in $\text{Time}(n^2)$. We can show that a substantial improvement would disprove SETH.

Theorem 4.7. $\text{Or-Vector} \in \text{SubquadraticTime} \Rightarrow \text{SETH}$ is false.

Proof. Divide the v variables in two blocks of $v/2$ each. For each assignment to the variables in the first block construct the vector in $[2]^{v/2}$ where bit i is 1 iff clause i is satisfied by the variables in the first block. Call A the resulting set of vectors. Let $N := 2^{v/2}$ and note $|A| = N$. Do the same for the other block and call the resulting set B .

Note that ϕ is satisfiable iff $\exists a \in A, b \in B$ such that $a \vee b = 1^d$.

Constructing these sets takes time Nd^c . If $\text{Or-Vector} \in \text{Time}(n^{2-\epsilon})$ for some $\epsilon > 0$, we can take $k = c_\epsilon$ and rule out SETH. **QED**

4.5 Search problems

Most of the problems in the previous sections ask about the *existence* of solutions. For example 3Sat asks about the existence of a satisfying assignment. It is natural to ask about

computing such a solution, if it exists. Such non-boolean problems are known as *search problems*.

Next we show that in some cases we can reduce a search problem to the corresponding boolean problem.

Definition 4.14. Search-3Sat is the problem: Given a satisfiable 3CNF formula, output a satisfying assignment.

Theorem 4.8. Search-3Sat reduces to 3Sat in P. That is: $3\text{Sat} \in \text{P} \Rightarrow \text{Search-3Sat} \in \text{P}$.

Proof. We construct a satisfying assignment one variable at the time. Given a satisfiable 3CNF, set the first variable to 0 and check if it is still satisfiable with the assumed algorithm for 3Sat. If it is, go to the next variable. If it is not, set the first variable to 1 and go to the next variable. **QED**

Exercise 4.7. Show $\text{Clique} \in \text{P} \Rightarrow \text{Search-Clique} \in \text{P}$.

4.5.1 Fastest algorithm for Search-3Sat

A curious fact about many search problems is that we know of an algorithm which is, in an asymptotic sense to be discussed now, essentially the fastest possible algorithm. This algorithm proceeds by simulating every possible program. When a program stops and outputs the answer, we can *check it* efficiently. Naturally, we can't just take any program and simulate it until it ends, since it may never end. So we will clock programs, and stop them if they take too long. There is a particular simulation schedule which leads to efficient running times.

Theorem 4.9. There is a RAM U such that on input any satisfiable formula x :

- (1) U outputs a satisfying assignment, and
- (2) If there is a RAM M that on input x outputs a satisfying assignment for x in t steps then U stops in $c_M t + |x|^c$ steps.

We are taking advantage of the RAM model. On other models it is not known if the dependence on t can be linear.

Proof. For $i = 1, 2, \dots$ the RAM U simulates RAM i for 2^i steps. Lemma 2.2 guarantees that for each i the simulation takes time $c2^i$. If RAM i stops and outputs y , then U checks in time $|x|^c$ if y is a satisfying assignment. If it is, then U outputs y and stops. Otherwise it continues.

Now let M be as in (2). As before, we work with an enumeration of programs where each program appears infinitely often. Hence we can assume that M has a description of length $\ell := c_M + \log t$. Thus the simulation will terminate when $i = \ell$.

The time spent by U for a fixed i is $\leq c \cdot 2^i + |x|^c$. Hence the total running time of U is

$$\leq c \sum_{j=1}^{\ell} (c2^j + |x|^c) \leq c_M 2^{\ell} + c_M |x|^c \leq c_M (t + |x|^c).$$

QED

This result nicely illustrates how “constant factors” can lead to impractical results because, of course, the problem is that the constant in front of t is enormous. Specifically, it is exponential in the size of the program, see Problem 4.6.

4.6 Gap-SAT: The PCP theorem

“Furthermore, most problem reductions do not create or preserve such gaps. There would appear to be a last resort, namely to *create* such a gap in the generic reduction [C]. Unfortunately, this also seems doubtful. The intuitive reason is that computation is an inherently unstable, non-robust mathematical object, in the sense that it can be turned from non-accepting by changes that would be insignificant in any reasonable metric – say, by flipping a single state to accepting.”

One of the most exciting, consequential, and technical developments in complexity theory of the last few decades has been the development of reductions that create *gaps*.

Definition 4.15. γ -Gap-3Sat is the 3Sat problem restricted to input formulas f that are either satisfiable or such that any assignment satisfies at most a γ fraction of clauses.

Note that 3Sat is equivalent to γ -Gap-3Sat for $\gamma = 1 - 1/n$, since a formula of size n has at most n clauses. At first sight it is unclear how to connect the problems when γ is much smaller. But in fact it is possible to obtain a constant γ . This result is known as the PCP theorem, where PCP stands for probabilistically-checkable-proofs. The connection to proofs will be discussed in Chapter 11.

Theorem 4.10. [PCP] There is $\gamma < 1$ such that γ -Gap-3Sat \in P \Rightarrow 3Sat \in P.

Similar results can be established for other problems such as 3Color, but the reductions in the previous section don’t preserve gaps and can’t be immediately applied.

A major application of the PCP theorem is in *inapproximability* results. A typical optimization problem is Max-3Sat.

Definition 4.16. The Max-3Sat problem: given a 3CNF formula, find a satisfying assignment that satisfies the maximum number of clauses.

Solving 3Sat reduces to Max-3Sat (in Chapter 5 we will give a reverse reduction as well). But we can ask for β -approximating Max-3Sat, that is, computing an assignment that satisfies a number of clauses that is at least a β fraction of the maximum possible clauses that can be satisfied.

The PCP Theorem 4.10 implies that 3Sat reduces to β -approximating Max-3Sat, for some constant $\beta < 1$.

It has been a major line of research to obtain tight approximation factors β for a variety of problems. For example, 3Sat reduces to β -approximating Max-3Sat for any $\beta > 7/8$. This constant is tight because a random uniform assignment to the variables satisfies each clause with probability $7/8$ and hence expects to satisfy a $7/8$ fraction of the clauses.

Exercise 4.8. Turn this latter observation in an efficient randomized algorithm with an approximation factor $7/8 - o(1)$.

4.7 Problems

Problem 4.1. Reduce 3Sat in P to the PIT problem (Definition 2.8) over the field with two elements.

Problem 4.2. Prove that 3Sat is not TM-Time($n^{1.99}$). (Hint: Consider a variant of the palindromes problem where the input bits are suitably spaced out with zeroes. Prove a time lower bound for this variant by explaining what modifications are needed to the proof of Theorem 3.1. Conclude by giving a suitable reduction from Padded-Palindromes to 3Sat.)

Problem 4.3. Consider the problem H : The input is a directed graph with a special source node s , m destination nodes t_1, t_2, \dots, t_m , and a subset B of *bad nodes*. The question is whether there are m paths from s to each of the destination nodes. The paths can share edges, but any two paths entering a bad node must leave through the same outgoing edge.

Reduce 3SAT to H in P.

Problem 4.4. Show that $3\text{Color} \in \text{P} \Rightarrow \text{Search-3Color} \in \text{P}$.

Problem 4.5. Give an encoding of 3Sat so that the reduction to 3Color in section §4.3 can be computed, for any input length, by a 1-local map (in particular, a circuit of constant depth) (cf. Definition 4.3).

Problem 4.6. Suppose there exists a such that Theorem 4.9 holds with the running time of U replaced with $(|M| \cdot t \cdot |x|)^a$. (That is, the dependence on the program description improved to power, and we allow even weaker dependence on t .) Prove that $3\text{Sat} \in \text{P}$.

Problem 4.7. Use Problem 2.6 and its notation. Assume $\text{P} = \text{BPP}$ (recall this only refers to boolean functions, see the convention on page 18). Show that given a circuit C and ϵ written in unary s.t. $p_C \geq \epsilon$ we can compute $x : C(x) = 1$ in P. In particular, given a non-zero arithmetic circuit we can find a non-zero assignment.

4.8 Notes

Circuits of size $cn \log n$ for multiplication were obtained in [83]. For the result about RAMs see [169].

Following [51] (discussed in the next chapter), [109] established reductions from satisfiability of (general) boolean formulas to 21 problems, including 3Sat, Clique, Cover-by-vertexes, 3Color, and Subset Sum. This opened the floodgates: The web of reductions from 3Sat is immense, see [64] for a starter.

The ETH and the SETH are from [99] and [101]. Again, a large number of reductions involving these hypotheses exists.

The web of reductions of 3Sum, including Theorem 4.2, was first spun in [62] and has grown ever since. Theorem 4.3 is from [209].

Theorem 4.9 is from [121].

The quote at the beginning of section §4.6 is from [149]. The PCP theorem as stated in Theorem 4.10 is from [17]. A sequence of exciting works preceded and followed it. For an account, as well as a proof of the PCP theorem, see [16].

Problem 4.7 is from [69].

Problem 4.6 is from [192].

This cluster of problems equivalent to 3Sat is so prominent that problems in it have been compiled into books [64], see also the list on wikipedia: https://en.wikipedia.org/wiki/List_of_NP-complete_problems. Amusingly, this list contains (generalized versions of) several popular games including: Tetris, Lemmings, Sudoku, etc. For an excellent exposition of this type of results see the video <https://www.youtube.com/watch?v=oS8m9fSk-Wk>

Tight hardness results based on SETH have been established for several well-studied problems, including longest-common subsequence [2] and edit distance [23].

Chapter 5

Completeness: Reducing arbitrary computation

In this chapter we show how to reduce arbitrary computation to 3Sat (and hence to the other problems in section §4.3). What powers everything is the following landmark and, in hindsight, simple result which reduces circuit computation to 3Sat.

Theorem 5.1. Given a circuit $C : [2]^n \rightarrow [2]$ with s gates we can compute in P a 3CNF formula f_C in $n + s$ variables such that for every $x \in [2]^n$:

$$C(x) = 1 \Leftrightarrow \exists y \in [2]^s : f_C(x, y) = 1.$$

The key idea to *guess computation and check it efficiently, using that computation is local*. The additional s variables one introduces contain the values of the gates during the computation of C on x . We simply have to check that they all correspond to a valid computation, and this can be written as 3CNF because each gate depends on at most two other gates.

Proof. Introduce a variable y_i for each non-input gate g_i in C . The value of y_i is intended to be the value of gate g_i during the computation. Whether the value of a gate g_i is correct is a function of 3 variables: y_i and the ≤ 2 gates that input g_i , some of which could be input variables. This can be written as a 3CNF by Theorem 2.4. Take an And of all these 3CNFs. Finally, add clause y_o for the output gate g_o . **QED**

Exercise 5.1. Write down the 3CNF for the circuit in figure 2.2, as given by the proof of Theorem 5.1.

Theorem 5.1 is a *depth-reduction* result. Indeed, note that a 3CNF can be written as a circuit of depth $c \log s$, whereas the original circuit may have any depth. This is helpful for example if you don't have the depth to run the circuit yourself. You can let someone else produce the computation, and you can check it in small depth.

We can combine Theorem 5.1 with the simulations in Chapter 2 to reduce computation in other models to 3SAT. In particular, we can reduce MTMs running in time t to 3Sat

of size $t \log^c t$. To obtain such parameters we need the quasilinear simulation of MTMs by circuits, Theorem 2.6.

However, recall that a quasilinear simulation of RAMs by circuits is not known. Only a power simulation is (which is obtained by combining the power simulation of RAMs by MTMs, Theorem 2.7, with a simulation of MTMs by circuits). This would reduce RAM computation running in time t to 3CNFs of size t^c . We content ourselves with this power loss for the beginning of this chapter. Later in section §5.3 we will obtain a quasi-linear simulation using an enjoyable argument which also bypasses Theorem 2.6.

In fact, these simulations apply to a more general, *non-deterministic*, model of computation. We define this model next, and then present the simulation with power loss in 5.2.

5.1 Nondeterministic computation

In the concluding equation in Theorem 5.1 there is an \exists quantifier on the right-hand side, but there isn't one on the left, next to the circuit. However, because the simulation works for every input, we can “stick” a quantifier on the left and have the same result. The resulting circuit computation $C(x, y)$ has two inputs, x and y . We can think of it as a *non-deterministic* circuit, which on input x outputs 1 iff $\exists y : C(x, y)$. Following the discussion before, we could do the same for other models like TMs, MTMs, and RAMs. The message here is that – if we allow for an \exists quantifier, or in other words consider nondeterministic computation – efficient computation is *equivalent* to 3CNF! This is one motivation for formally introducing a *nondeterministic* computational model.

Definition 5.1. $\text{NTime}(t(n))$ is the set of functions $f : X \subseteq [2]^* \rightarrow [2]$ for which there is a RAM M such that:

- $f(x) = 1$ iff $\exists y \in [2]^{t(|x|)}$ such that $M(x, y) = 1$, and
- $M(x, y)$ stops within $t(|x|)$ steps on every input (x, y) .

We also define

$$\text{NP} := \bigcup_{d \geq 1} \text{NTime}(n^d),$$

$$\text{NExp} := \bigcup_{d \geq 1} \text{NTime}(2^{n^d}).$$

Note that the running time of M is a function of $|x|$, not $|(x, y)|$. This difference is inconsequential for NP, since the composition of two powers is another power. But it is important for a more fine-grained analysis. We refer to a RAM machine as in Definition 5.1 as a *nondeterministic machine*, and to the y in $M(x, y)$ as the *nondeterministic choices*, or *guesses*, of the machine on input x .

We can also define NTime in a way that is similar to BPTIME, Definition 2.7. The two definitions are essentially equivalent. Our choice for BPTIME is motivated by the identification of BPTIME with computation that is actually run. For example, in a programming

language one uses an instruction like Rand to obtain random values; one does not think of the randomness as being part of the input. By contrast, NTime is a more abstract model, and the definition with the nondeterministic guesses explicitly laid out is closer in spirit to a 3CNF.

All the problems we studied in section §4.3 are in NP.

Fact 5.1. 3Sat, Clique, Cover-by-vertexes, SubsetSum, and 3Color are in NP.

Proof. For a 3Sat instance f , the variables y correspond to an assignment. Checking if the assignment satisfies f is in P. This shows that 3Sat is in NP. **QED**

Exercise 5.2. Finish the proof by addressing the other problems in Fact 5.1

5.1.1 How to think of NP

We can think of NP as the problems which admit a solution that can be verified efficiently, namely in P. For example for 3Sat it is easy to verify if an assignment satisfies the clauses, for 3Color it is easy to verify if a coloring is such that any edge has endpoints of different colors, for SubsetSum it is easy to verify if a subset has a sum equal to a target, and so on. However, as we saw above this verification step can be cast in a restricted model, namely a 3CNF. So we don't have to think of the verification step as using the full power of P computation.

Here's a vivid illustration of NP. Suppose I claim that the following matrix contains a 9:

56788565634705634705637480563476
70156137805167840132838202386421
85720582340570372307580234576423
80275880237505788075075802346518
78502378564067807582348057285428
05723748754543650350562378804337
52305723485008160234723884077764
86543234567865435674567836738063
45463788486754345743457483460040
73273873486574375464584895741832
85075783485634856237847287422112
83748874883753485745788788223201

How can you tell, without tediously examining the whole matrix? However, if I tell you that it's in row 10, 8 digits from the right, you can quickly check that I am right. I won't be able to cheat, since you can check my claims. On the other hand I can provide a proof that's easy to verify.

P vs. NP

The flagship question of complexity theory is whether $P = NP$ or not. This is a young, prominent special case of the grand challenge we introduced in Chapter 3. Contrary to the analogous question for BPP, cf. section 2.6.2, the general belief seems to be that $P \neq NP$. Similarly to BPP, cf. Theorem 2.10, the best deterministic simulation of NP runs in exponential time by trying all nondeterministic guesses. This gives the middle inclusion in the following fact; the other two are by definition.

Fact 5.2. $P \subseteq NP \subseteq Exp \subseteq NExp$.

A consequence of the Time Hierarchy Theorem 3.4 is that $P \neq Exp$. From the inclusions above it follows that

$$P \neq NP \text{ or } NP \neq Exp, \text{ possibly both.}$$

Thus, we are not completely clueless, and we know that at least one important separation is lurking somewhere. Most people appear to think that *both* separations hold, but we are unable to prove *either*.

For multi-tape machines, a separation between deterministic and non-deterministic linear time is in [152, 166].

5.2 NP-completeness

We now go back to the question at the beginning of this chapter about reducing arbitrary computation to 3Sat. We shall reduce all of NP to 3Sat in Theorem 5.2. Problems like 3Sat admitting such reductions deserve a definition.

Definition 5.2. We call a problem L :

NP-hard if every problem in NP reduces to L in P;

NP-complete if it is NP-hard and in NP. To spell it out, this means that $L \in NP$ and moreover for any $M \in NP$ we have $L \in P \Rightarrow M \in P$.

One can define NP-hard (and hence NP-complete) w.r.t. different reductions, cf. Chapter 4, and we will do so later. But the simple choice above suffices for now.

Complete problems are the “hardest problems” in the class, as formalized in the following fact.

Fact 5.3. Suppose L is NP-complete. Then $L \in P \Leftrightarrow P = NP$.

Proof. (\Leftarrow) This is because $L \in NP$.

(\Rightarrow) Let $L' \in NP$. Because L is NP-hard we know that $L \in P \Rightarrow L' \in P$. **QED**

Exercise 5.3. Suppose $P = NP$. Prove that any problem in NP is NP-complete.

Suppose instead $P \neq NP$. Let $L \in NP$. Prove L is NP-complete iff $L \notin P$.

Fact 5.3 points to an important interplay between problems and complexity classes. We can study complexity classes by studying their complete problems, and vice versa.

The central result in the theory of NP completeness is the following.

Theorem 5.2. 3Sat is NP-complete.

Proof. 3Sat is in NP by Fact 5.1. Next we prove NP-hardness. The main idea is Theorem 5.1, while the rest of the proof mostly amounts to opening up definitions and using some

previous simulations. Let $L \in \text{NP}$ and let M be the corresponding TM which runs in time n^d on inputs (x, y) where $|x| = n$ and $|y| = n^d$, for some constant d . We can work with TMs instead of RAMs since they are equivalent up to a power loss, as we saw in Theorem 2.7. We can construct in P a circuit $C(x, y)$ of size $c_M n^{c_d}$ such that for any x, y we have $M(x, y) = 1 \Leftrightarrow C(x, y) = 1$ by Theorem 2.5.

Now, suppose we are given an input w for which we are trying to decide membership in L . This is equivalent to deciding if $\exists y : C(w, y) = 1$ by what we just said. We can “hard-wire” w into C to obtain the circuit $C_w(y) := C(w, y)$ only on the variables y , with no loss in size. Here by “hard-wire” we mean replacing the input gates x with the bits of w . Now we can apply Theorem 5.1 to this new circuit to produce a 3CNF f_w on variables y and new variables z such that $C_w(y) = 1 \Leftrightarrow \exists z : f_w(y, z) = 1$, for any y . The size of f_w and the number of variables z is power in the size of the circuit.

We have obtained:

$$w \in L \Leftrightarrow \exists y : M(w, y) = 1 \Leftrightarrow \exists y : C_w(y) = 1 \Leftrightarrow \exists y, z : f_w(y, z) = 1 \Leftrightarrow f_w \in \text{3Sat},$$

as desired. **QED**

In section §4.3 we reduced 3Sat to other problems which are also in NP by Fact 5.1. This implies that all these problems are NP-complete. Here we use that if problem A reduces to B in P , and B reduces to C , then also A reduces to C . This is because if $C \in \text{P}$ then $B \in \text{P}$, and so $A \in \text{P}$.

Corollary 5.1. Clique, Cover-by-vertexes, Subset-sum, and 3Color are NP-complete.

It is important to note that there is nothing special about the *existence* of NP-complete problems. The following is a simple such problem that does not require any of the machinery in this section.

Exercise 5.4. Consider the problem, given a RAM M , an input x , and $t \in \mathbb{N}$, where t is written in unary, decide if there is $y \in [2]^t$ such that $M(x, y) = 1$ in t steps. Prove that this is NP-complete.

What if t is written in binary?

The interesting aspect of NP-complete problems such as 3Sat and those in Corollary 5.1 is that they are very simple and structured, and don’t refer to computational models. This makes them suitable for reductions, and for inferring properties of the complexity class which are not evident from a machine-based definition.

5.3 From RAM to 3SAT in quasi-linear time

The framework in the previous section is useful to relate membership in P of different problems in NP, but it is not suitable for a more fine-grained analysis. For example, under the assumption that 3Sat is in $\text{Time}(cn)$ we cannot immediately conclude that other problems in NP are solvable in this time or in about this time. We can only conclude that they are in P .

In particular, the complexity of 3Sat cannot be related to that of other central conjectures, such as whether 3Sum is in subquadratic time, Conjecture 4.1.

The culprit is the power loss in reducing RAM computation to circuits, mentioned at the beginning of the chapter. We now remedy this situation and present a quasi-linear reduction. As we did before, cf. Theorem 5.1 and Theorem 5.2, we first state a version of the simulation for (deterministic) computation which contains all the main ideas, and then we note that a completeness result follows.

Theorem 5.3. Given an input length $n \in \mathbb{N}$, a time bound $t \in \mathbb{N}$, and a RAM M that runs in time t on inputs of n bits, we can compute in time $t' := c_M t (\log t)^c$ a 3CNF f on variables (x, y) where $|y| \leq t'$ such that for every $x \in [2]^n$:

$$M(x) = 1 \iff \exists y : f(x, y) = 1.$$

We now present the proof of this amazing result; you may want to refer back to Definition 2.5 of a RAM. A key concept in the proof is the following “snapshot” of the RAM computation.

Definition 5.3. The *internal configuration*, abbreviated IC, of a RAM specifies:

- its registers,
- the program counter,
- the word length w , and
- if the current instruction is a Read $r_i := \mu[r_j]$ or Write $\mu[r_j] := r_i$ then the IC includes the content $\mu[r_j]$ of the memory cell indexed by r_j .

Note that at most one memory cell is included in one IC. By contrast, the configuration of a TM (Definition 2.1) includes all its tape cells. Also note that an IC has length $\leq c_M + c \log t$ bits, where the c_M is for the program counter, and the $c \log t$ is for the rest, using that the maximum word length of a machine running in time $t \geq n$ is $c \log t$.

The key idea in the proof. At the high level, the approach is, like in Theorem 5.1, to guess computation and check it efficiently. We are going to *guess* the sequence of ICs, and we need additional ideas to check them efficiently by a circuit. This is not immediate, since, again, the RAM can use direct access to read and write in memory at arbitrary locations, something which is not easy to do with a circuit.

The key idea is to check operations involving memory *independently* from the operations involving registers but not memory. If both checks pass, then the computation is correct. More precisely, a sequence of internal configurations s_1, s_2, \dots, s_t corresponds to the computation of the RAM on input x iff for every $i < t$:

1. If s_i does not access memory, then s_{i+1} has its registers, program counter, and word length updated according to the instruction executed in s_i ,

2. If s_i is computing a read operation $r_i := \mu[r_j]$ then in s_{i+1} register r_j contains *the most recent value written in memory cell r_j* . In case this cell was never written, then r_j should contain x_j if $j \in \{1, 2, \dots, n\}$, n if $j = 0$, and 0 otherwise. The program counter in s_{i+1} also points to the next instruction.

Rather than directly constructing a 3CNF that implements these checks, we construct a circuit and then appeal to Theorem 5.1. It is easy to construct a circuit of quasi-linear size implementing Check 1, since the circuit only has to check adjacent pairs of ICs. As remarked before, these ICs have length $\leq c_M + c \log t$. For fixed i , Check 1 can be implemented by a circuit which depends on the RAM and has size power in the length of an IC. Taking an And of these circuits over the choices of i gives a circuit of the desired size for Check 1.

The difficulty lies in Check 2, because the circuit needs to find “the most recent value written.” The solution is to *sort* the ICs by memory addresses. After sorting, we can implement Check (2) as easily as Check (1), since we just need to check adjacent pairs of ICs.

The emergence of sorting in the theory of NP-completeness cements the pivotal role this operation plays in computer science.

To implement this idea we need to be able to sort with a quasi-linear size circuit. Standard sorting algorithms like Mergesort, Heapsort, or Radixsort run in quasi-linear time on a RAM, but rely on direct addressing (cf. section §2.5) and for this reason cannot be easily implemented by a circuit of quasi-linear size. However other algorithms have been developed that do have such an implementation. This gives the following lemma.

Lemma 5.1. Given t and m we can compute in time $t' := t \cdot (m \log t)^c$ a circuit (of size $\leq t'$) that sorts t integers of m bits.

Because this reduction is so fundamental, for completeness we give a proof of Lemma 5.1 in section §5.3.1.

We summarize the key steps in the proof.

Proof of Theorem 5.3. We construct a circuit C_M and then appeal to Theorem 5.1. The extra variables y correspond to t ICs s_1, s_2, \dots, s_t . An IC takes $c_M + c \log t$ bits to specify, so we need $\leq c_M t \log t$ variables y . The circuit C_M first performs Check (1) above for each adjacent pair (s_i, s_{i+1}) of ICs. This takes size $c_M \log^c t$ for each pair, and so size $c_M t \log^c t$ overall.

Then C_M sorts the ICs by memory addresses, producing sorted ICs s'_1, s'_2, \dots, s'_t . This takes size $t \cdot \log^c t$ by Lemma 5.1, using that the memory addresses have $\leq c \log t$ bits. Then the circuit performs Check (2) for each adjacent pair (s'_i, s'_{i+1}) of ICs. The circuit size required for this is no more than for Check (1).

Finally, the circuit takes an And of the results of the two checks, and also checks that s_t is accepting. **QED**

We can now prove completeness in a manner similar to Theorem 5.2, with a relatively simple extension of Theorem 5.3.

Theorem 5.4. Every problem L in $\text{NTime}(t)$ map reduces to 3Sat in $\text{Time}(c_{L,t}t \log^c t)$, for every function $t \geq n$ such that $t(x)$ is computable in time $t(x)$ given x .

The assumption on t is similar to that in the hierarchy Theorem 3.4, and is satisfied by all standard functions including all those in this book – cf. discussion after Theorem 3.4.

Proof. Let M be a RAM computing L in the assumed time. Given an input w of length n we have to efficiently compute a 3CNF f such that

$$\exists y \in [2]^{t(n)} : M(w, y) = 1 \iff \exists y \in [2]^{c_{L,t}t \log^c t(n)} : f(y) = 1.$$

First we compute $t(n)$, using the assumption. We now apply Theorem 5.3, but on a new input length $n' := c(n+t) \leq ct$, to accommodate for inputs of the form (x, y) . This produces a formula f of size $c_{L,t}t(\log t)^c$ in variables (x, y) and new variables z . We can now set x to w and conclude the proof. **QED**

With these sharper results we can now study hardness and completeness within time bounds such as n^2 , $n \log^3 n$ etc. We work out an example in the next section.

5.3.1 Efficient sorting circuits: Proof of Lemma 5.1

We present an efficient sorting algorithm for an array $A[n]$ which enjoys the following property: *the only way in which the input is accessed is via Compare-Exchange operations.* Compare-Exchange takes two indexes i and j and swaps $A[i]$ and $A[j]$ if they are in the wrong order. It has the following code:

```
Compare-Exchange(Array  $A[0..(n-1)]$  and indexes  $i$  and  $j$  with  $i < j$ ):
  if  $A[i] > A[j]$ 
    swap  $A[i]$  and  $A[j]$ 
```

Why care about this property? It makes the comparisons *independent from the data*, and this allows us to implement the algorithm with a network – a *sorting network* – of fixed Compare-Exchange operations. In particular, we will get a circuit.

We call an algorithm with this property *oblivious*. Familiar mergesort is not oblivious, because the merge operations performs comparisons which depend on the outcome of previous ones. However a variant of Mergesort [26], called Odd-Even-Mergesort, is oblivious.

Algorithm Odd-Even-Merge(A) merges the two already sorted halves $[a_0, a_1, \dots, a_{n/2-1}]$ and $[a_{n/2}, a_{n/2+1}, \dots, a_{n-1}]$ of the sequence $A = [a_0, a_1, \dots, a_{n-1}]$, resulting in a sorted output sequence. It works in a remarkable and mysterious way. First it merges the *odd* subsequence of the entire array A , then the *even*, and finally it makes $O(n)$ Compare-Exchange-Operations. Throughout, we assume that n is a power of 2.

```
Odd-Even-Merge( $A = [a_0, \dots, a_{(n-1)}]$ ):
  if  $n = 2$ 
    Compare-Exchange( $A, 0, 1$ )
```



```

else {
  Odd-Even-Merge( $[a_0, a_2, \dots, a_{(n-2)}]$ ,  $n/2$ ) //the even subsequence
  Odd-Even-Merge( $[a_1, a_3, \dots, a_{(n-1)}]$ ,  $n/2$ ) //the odd subsequence
  for  $i \in \{1, 3, 5, 7, \dots, n-3\}$ 
    Compare-Exchange( $A, i, i+1$ )
}

```

We shall now argue that this algorithm is correct.

Lemma 5.2. If $[a_0, a_1, \dots, a_{n/2-1}]$ and $[a_{n/2}, a_{n/2+1}, \dots, a_{n-1}]$ are sorted, then Odd-Even-Mergesort($[a_0, a_1, \dots, a_{n-1}]$) outputs a sorted array.

Proof. To prove this lemma we invoke the so-called “0-1 principle.” This principle says that it suffices to prove the lemma when each a_i is either 0 or 1, assuming that the algorithm only accesses the input via Compare-Exchange operations. For completeness we sketch a proof of this principle in this paragraph. Let $A = [a_0, \dots, a_{n-1}]$ be an input to Odd-Even-Merge, and let $B = [b_0, \dots, b_{n-1}]$ be the output sequence produced by the algorithm. If the algorithm fails to correctly sort A , then consider the smallest index k such that $b_k > b_{k+1}$. Define a function f such that $f(c) = 1$ if $c \geq b_k$ and $f(c) = 0$ otherwise. For an array $X = [X_0, X_1, \dots, X_{n-1}]$ let $f(X)$ be the sequence $[f(X_0), f(X_1), \dots, f(X_{n-1})]$ obtained by applying f to each element of X . Observe that $f(B)$ is not sorted. However it is easy to verify that f commutes with any Compare-Exchange operation applied to any sequence X , i.e.,

$$f(\text{Compare-Exchange}(X, i, j)) = \text{Compare-Exchange}(f(X), i, j).$$

Because Odd-Even-Merge is just a sequence of Compare-Exchange, we have that


$$f(B) = f(\text{Odd-Even-Merge}(A)) = \text{Odd-Even-Merge}(f(A))$$

and so the algorithm fails to correctly merge the 0-1 sequence $f(A)$. It only remains to notice that $f(A)$ is a valid input for Odd-Even-Merge. This is indeed the case because if a sequence X is sorted then $f(X)$ is also sorted.

We now prove the lemma by induction on n , based on the recursive definition of Odd-Even-Merge. Refer to Figure 5.3.1.

The base case $n = 2$ is clear. Assume that Odd-Even-Merge correctly merges any two sorted 0-1 sequences of size $n/2$. We view an input sequence of n elements as an $n/2 \times 2$ matrix, with the left column corresponding to elements at the even-indexed positions $0, 2, \dots, n-2$ and the right column corresponding to elements at the odd-indexed positions $1, 3, \dots, n-1$ (Figure 5.3.1(a)). Figure 5.3.1(b) shows a corresponding 0-1 input, which we can assume w.l.o.g. because of the zero-one principle. Figure 5.3.1(c) shows the matrix after the recursive calls to the sorting. Since the upper half of the matrix is sorted by assumption, the right column in the upper half has the same number or exactly one more 1 than the left column in the upper half. The same is true for the lower half. Because each (length- $(n/4)$) column in each half of the matrix is also individually sorted by assumption, the induction

a[0]	a[1]
a[2]	a[3]
a[4]	a[5]
a[6]	a[7]
a[8]	a[9]
a[10]	a[11]
a[12]	a[13]
a[14]	a[15]



1

Figure 5.1: Odd-Even-Mergesort

hypothesis guarantees that after the two calls to Odd-Even-Merge both the left and right (length- $(n/2)$) columns are sorted (Figure 5.3.1(d)).

At this point only one of 3 cases arises:

- 1) The odd and even subsequences have the same number of 1s.
- 2) The odd subsequence has a single 1 more than the even subsequence.
- 3) The odd subsequence has two 1s more than the even subsequence.

In the first two cases, the sequence is already sorted. In the third case, the Compare-Exchange operations (Figure 5.3.1(e)) yield a sorted sequence (Figure 5.3.1(f)). **QED**

Given Odd-Even-Merge, we can sort by the following algorithm which has the same structure as Mergesort

```
Oblivious-Mergesort( $A = [a_0, \dots, a_{n-1}]$ ):
if  $n \geq 2$  {
  Oblivious-Mergesort( $[a_0, a_1, \dots, a_{n/2-1}]$ )
  Oblivious-Mergesort( $[a_{n/2}, a_{n/2+1}, \dots, a_{n-1}]$ )
  Odd-Even-Merge( $[a_0, a_1, \dots, a_{n-1}]$ )
}
```

It only remains to argue efficiency. Let $S_M(n)$ denote the number of Compare-Exchange operations for Odd-Even-Merge for an input sequence of length n . We have the recurrence

$$S_M(n) = 2 \cdot S_M(n/2) + n/2 - 1,$$

which yields $S_M(n) = O(n \cdot \log n)$.

Finally, let $S(n)$ denote the number of calls to Compare-Exchange for Oblivious-Mergesort with an input sequence of length n . Then we have the recurrence $S(n) = 2 \cdot S(n/2) + (n \cdot \log n)$, which yields $S(n) = O(n \cdot \log^2 n)$.

To conclude the proof, note that Compare-Exchange for inputs with m bits can be implemented by a circuit of size m^c . **QED**

5.3.2 Quasilinear-time completeness

In this section we use the machinery we just developed to study completeness in quasi-linear time, instead of power time.

Definition 5.4. We define the quasi-linear time complexity classes

$$\begin{aligned} \text{QLin-Time} &:= \bigcup_{d \in \mathbb{N}} \text{Time}(n \log^d n) \text{ and} \\ \text{QLin-NTime} &:= \bigcup_{d \in \mathbb{N}} \text{NTime}(n \log^d n). \end{aligned}$$

Theorem 5.5. 3Sat is complete for QLin-NTime with respect to mapping reductions in QLin-Time. That is:

- 3Sat is in QLin-NTime, and
- every problem in QLin-NTime map reduces to 3Sat in QLin-Time.

Proof. To show that 3Sat is in QLin-NTime, consider a 3CNF instance f of length n . This instance has at most n variables, and we can guess an assignment y to them within our budget of non-deterministic guesses. There remains to verify that y satisfies f . For this, we can do one pass over the clauses. For each clause, we access the bits in y corresponding to the 3 variables in the clause, and check if the clause is satisfied. This takes constant time per clause, and so time cn overall.

The second part follows from Theorem 5.4, using the fact that the composition of two quasilinear functions is also quasilinear (similarly to the fact that the composition of two power functions is also a power). **QED**

Note that the proof that 3Sat is in QLin-NTime relies on our computational model being a RAM, because we use direct access to fetch the values for the variables in a clause.

We can now give the following quasi-linear version of Fact 5.3. The only extra observation for the proof is again that the composition of two quasi-linear functions is quasi-linear.

Corollary 5.2. $3\text{Sat} \in \text{QLin-Time} \Leftrightarrow \text{QLin-NTime} = \text{QLin-Time}$.

Exercise 5.5. Prove that Theorem 5.5 holds with 3Color instead of 3Sat. What about Clique and Subset-sum?

Exercise 5.6. Prove that 3Sum reduces to 3Sat in Subquadratic time. That is: $3\text{Sat} \in \text{SubquadraticTime} \Rightarrow 3\text{Sum} \in \text{SubquadraticTime}$ (i.e., Conjecture 4.1 is false).

5.4 Completeness in other classes

The completeness phenomenon is not special to NP but enjoyed by many other classes. In this section we begin to explore completeness for NExp and Exp. One needs to be careful how hardness (and hence completeness) is defined, since these classes are known to be different from P by the hierarchy Theorem 3.4. So defining a problem L to be NExp-hard if $L \in \text{P} \Rightarrow \text{NExp} = \text{P}$ would mean simply that $L \notin \text{P}$. To avoid this in this section hardness (hence completeness) is defined w.r.t. mapping reductions, cf. Chapter 4. (Another option would be to replace P with say BPP, since it is not known if $\text{BPP} = \text{NExp}$.)

5.4.1 NExp completeness

Complete problems for NExp include *succinct* versions of problems complete for NP. Here succinct means that rather than giving the input x to the problem in standard format, the input consists instead of a circuit $C : [2]^m \rightarrow [2]$ encoding x , for example $C(i)$ equals bit i of x , for every i .

Definition 5.5. The Succinct-3Sat problem: Given a circuit C encoding a 3CNF f_C , does f_C have a satisfying assignment?

Theorem 5.6. Succinct-3Sat is NExp complete with respect to power-time mapping reductions.

Proof sketch. Let us first show that Succinct-3Sat is in NExp. Given a circuit C of length n , we can run it on every possible input (of length $\leq n$) and write down the formula f_C encoded by C . This formula has size $\leq 2^n$. We can then use the fact that 3Sat is in NP to decide satisfiability of this formula in non-deterministic power time in 2^n , that is $\text{NTime}(2^{cn}) \subseteq \text{NExp}$.

To prove NExp hardness it is convenient to work with TMs rather than RAMs. The main observation is that in the simulation of a TM M on an input x by a circuit C_M , Theorem 2.5, the circuit is very regular, in the sense that we can construct another circuit S_M which is a succinct encoding of C_M . The circuit S_M is given as input indexes to gates in C_M and outputs the type of the gate and its wires. The size of S_M is power in the index length and M . Thus, if C_M has size t^c , S_M only needs size $\log^c t$. If $t = 2^{n^d}$, S_M has size power in n , as desired. The transformation from circuit to 3CNF in Theorem 5.1 is also regular and can be done succinctly. **QED**

As a consequence, we obtain the following “concrete” problem not in P.

Corollary 5.3. Succinct-3Sat \notin P.

5.4.2 Exp-completeness

Exp-complete problems include several two-player games. The important feature for completeness is that the game may last for an exponential number of steps (otherwise it would belong to a class believed to be stricter which we will investigate in Chapter 7). These games include (generalized versions of) Chess and Checkers.

5.5 Power from completeness

The realization that arbitrary computation can be reduced to 3Sat and other problems is powerful and liberating. In particular it allows us to significantly widen the net of reductions.

5.5.1 Optimization problems

As observed in section §4.6, 3Sat trivially reduces to Max-3Sat. The converse will be shown next.

Theorem 5.7. Max-3Sat reduces to 3Sat in P.

Proof. Consider the problem *Atleast-3Sat*: Given a 3CNF formula and an integer t , is there an assignment that satisfies at least t clauses? This is in NP and so can be reduced to 3Sat in P. This is the step that's not easy without "thinking completeness:" given an algorithm for 3Sat it isn't clear how to use it directly to solve *Atleast-3Sat*.

Hence, if 3Sat is in P so is *Atleast-3Sat*. On input a 3CNF f , using binary search and the fact that *Atleast-3Sat* is in P, we can find in P the largest t s.t. $(f, t) \in \text{Atleast-3Sat}$. Having found this t , there remains to construct an assignment satisfying the clauses. This can be done fixing one variable at the time as in Theorem 4.8. **QED**

5.5.2 NP is as easy as detecting unique solutions

A satisfiable 3CNF can have multiple satisfying assignments. On the other hand some problems and puzzles have unique solutions. In this section we relate these two scenarios.

Definition 5.6. *Unique-CktSat* is the problem: Given a circuit C s.t. there is at most one input x for which $C(x) = 1$, decide if such an input exists.

Unique-3Sat is the *Unique-CktSat* problem restricted to 3CNF circuits.

Theorem 5.8. 3Sat reduces to *Unique-3Sat* in BPP.

We in fact reduce 3Sat to *Unique-CktSat*. Then *Unique-CktSat* can be reduced to *Unique-3Sat* observing that the reduction in Theorem 5.1 preserves uniqueness.

The beautiful proof uses a powerful and general technique in randomized computation: *pairwise uniformity*, sometimes more generically referred to as *hashing*. We first define such functions and give efficient constructions. Then we show how to use them to "isolate" assignments.

Definition 5.7. A distribution H on functions mapping $S \rightarrow T$ is called *pairwise uniform* if for every $x, x' \in S$ and $y, y' \in T$ one has

$$\mathbb{P}_H[H(x) = y \wedge H(x') = y'] = 1/|T|^2.$$

This is saying that on every pair of inputs H is behaving as a completely uniform function. Yet unlike completely uniform functions, the next lemma shows that pairwise uniform functions can have a short description, which makes them suitable for use in algorithms.

Exercise 5.7. Let \mathbb{F}_q be a finite field. Define the random function $H : \mathbb{F}_q \rightarrow \mathbb{F}_q$ as $H(x) := Ax + B$ where A, B are uniform in \mathbb{F}_q .

Prove that H is pairwise uniform.

Explain how to use H to obtain a pairwise uniform function from $[2]^n$ to $[2]^t$ for any given $t \leq n$.

Exercise 5.8. Define the random function $H_1 : [2]^n \rightarrow [2]$ as $H_1(x) := \sum_{i \leq n} A_i x_i + B \pmod{2}$ where A is uniform in $[2]^n$ and B is uniform in $[2]$.

Prove that H_1 is pairwise uniform.

Explain how to use H_1 to obtain a pairwise uniform function from $[2]^n$ to $[2]^t$ for any given $t \leq n$.

We can now state the lemma that we use to isolate assignments.

Lemma 5.3. Let H be a pairwise uniform function mapping $S \rightarrow T$, and let $1 \in T$. The probability that there is a unique element $s \in S$ such that $H(s) = 1$ is

$$\geq \frac{|S|}{|T|} - \frac{|S|^2}{|T|^2}.$$

In particular, if $|T|/8 \leq |S| \leq |T|/4$ this prob. is $\geq \frac{1}{8} - \frac{1}{16} \geq 1/8$.

Proof. For fixed $s \in S$, the probability s is the unique element mapped to 1 is at least the prob. that s is mapped to 1 minus the prob. that both s and some other $s' \neq s$ are mapped to 1. This is

$$\geq \frac{1}{|T|} - \frac{|S| - 1}{|T|^2}.$$

These events for different $s \in S$ are disjoint; so the target probability is at least the sum of the above over $s \in S$. **QED**

Proof of Theorem 5.8. Given a 3Sat instance ϕ with $\leq n$ variables x , we pick a random i from 0 to $n + c$. We then pick a pairwise uniform function mapping $[2]^n$ to $[2]^i$, and consider the circuit

$$C := \phi(x) \wedge H(x) = 0^i.$$

This circuit has size n^c .

If ϕ is not satisfiable, C is not satisfiable, for any random choices.

Now suppose that ϕ has $s \geq 1$ satisfying assignment. With prob. $\geq 1/n$ we will have $2^{i-3} \leq s \leq 2^{i-2}$, in which case Lemma 5.3 guarantees that C has a unique satisfying assignment with prob. $\geq c$.

Overall, C has a unique satisfying assignment with prob. $\geq c/n$. Hence the Unique-3Sat algorithm on C outputs 1 with prob. $\geq c/n$. If we repeat this process cn times, with independent random choices, the Or of the outcomes gives the correct answer with prob. $\geq 2/3$. **QED**

5.6 Problems

Problem 5.1. In this problem you will explore an alternative proof of the results in section §5.3.

(1) Show $\text{Time}(t) \subseteq \exists \cdot c\text{-TM-Time}(t \log^c t)$. Note the lhs is for RAMs, the rhs is for TMs with c tapes. Hint: Follow the proof structure in section §5.3. Work with plain MergeSort. (In case you are unfamiliar with MergeSort, you might want to review it.)

(2) Prove Theorem 5.3 from (1) using a simulation from Chapter 2.

Problem 5.2. In Theorem 4.8 we reduced Search-3Sat to 3Sat.

(1) Suppose 3Sat is computable by circuits of depth $c \log n$. What would be the depth of the circuits for Search-3Sat given by the reduction?

(2) Reduce Search-3Sat to 3Sat in $\bigcup_{a>0} \text{Depth}(a \log n)$. Hint: First work with randomized circuits. Use ideas in the proof of 5.8. Then explain how to get deterministic circuits.

Note: Depending on how you feel about log-depth circuits, this problem belongs to either Chapter 5 or Chapter 9.

5.7 Notes

NP-completeness and Theorem 5.2 originates in the fundamental works [51, 121]. The first paper proves a version of Theorem 5.1 for TMs, for a more recent and similar exposition see [178]. Theorem 5.3 is from [78, 164]. The first work focuses on an equivalence between computational models, while the second explicitly constructs a 3CNF formula. We presented the proof in a slightly different way, using the sorting circuits from [26] and following the exposition in [142].

For the Exp-completeness of Chess see [59]; for Checkers [163].

The reduction to unique-3Sat, Theorem 5.8, is from [198] from which we borrowed the section title which, interestingly, emphasizes how easy NP could be, cf. Chapter 17.

Pairwise uniformity was studied as least since [155] and [44]. For background see [194, 87].

Chapter 6

Alternation

We placed one quantifier “in front” of computation and got something interesting: NP. So let’s push the envelope.

Definition 6.1. $\Sigma_i\text{Time}(t(n))$ is the set of functions $f : X \subseteq [2]^* \rightarrow [2]$ for which there is a RAM M such that on input $(x, y_1, y_2, \dots, y_i)$ stops within $t(|x|)$ steps and

$$f(x) = 1 \Leftrightarrow \exists y_1 \forall y_2 \exists y_3 \forall y_4 \dots Q_i y_i \in [2]^{t(|x|)} : M(x, y_1, y_2, \dots, y_i) = 1.$$

$\Pi_i\text{Time}(t(n))$ is defined similarly except that we start with a \forall quantifier. We also define

$$\Sigma_i\text{P} := \bigcup_d \Sigma_i\text{Time}(n^d),$$

$$\Pi_i\text{P} := \bigcup_d \Pi_i\text{Time}(n^d), \text{ and}$$

$$\text{the power hierarchy PH} := \bigcup_i \Sigma_i\text{P} = \bigcup_i \Pi_i\text{P}.$$

We refer to such computation and the corresponding machines as *alternating*, since they involve alternation of quantifiers and we will soon see a connection with alternating circuits.

The PH is the analog of the older arithmetical hierarchy from computability theory or logic in which nondeterministic time plays the role of listable (a.k.a. computably enumerable, etc.).

As was the case for NP, Definition 5.1, note that the running time of M is a function of $|x|$ only. Again, this difference is inconsequential for $\Sigma_i\text{P}$, since the composition of two powers is another power. But it is important for a more fine-grained analysis.

Exercise 6.1. Min-Ckt is the problem of deciding if an input circuit has an equivalent circuit which is smaller. It is not known to be in NP. In which of the above classes can you place it?

6.1 Does the PH collapse?

We refer to the event that $\exists i : \Sigma_i P = PH$ as “the PH collapses.” It is unknown if the PH collapses. Most people appear to believe that it does not, and to consider statements of the type

$$X \Rightarrow PH \text{ collapses}$$

as evidence that X is false. Examples of such statements are discussed next.

Theorem 6.1. $P = NP \Rightarrow P = PH$.

The idea in the proof is simply that if you can remove a quantifier then you can remove more.

Proof. We prove by induction on i that $\Sigma_i P \cup \Pi_i P = P$.

The base case $i = 1$ follows by assumption and the fact that P is closed under complement.

Next we do the induction step. We assume the conclusion is true for i and prove it for $i + 1$. We will show $\Sigma_{i+1} P = P$. The result about $\Pi_{i+1} P$ follows again by complementing.

Let $L \in \Sigma_{i+1} P$, so $\exists a$ and a power-time TM M such that for any $x \in [2]^n$,

$$x \in L \Leftrightarrow \exists y_1 \forall y_2 \exists y_3 \forall y_4 \dots Q_{i+1} y_{i+1} \in [2]^{n^a} : M(x, y_1, y_2, \dots, y_{i+1}) = 1.$$

(As discussed after Definition 6.1 we don't need to distinguish between time as a function of $|x|$ or of $|(x, y_1, y_2, \dots, y_{i+1})|$ when considering power times as we are doing now.)

Now the creative step of the proof is to consider

$$L' := \{(x, y_1) : \forall y_2 \in [2]^{n^a} \dots Q_{i+1} y_{i+1} \in [2]^{n^a} : M(x, y_1, y_2, \dots, y_{i+1}) = 1\}.$$

Note $L' \in \Pi_i P$. By induction hypothesis $L' \in P$. So let TM M' solve L' in power time. So $x \in L \iff \exists y_1 \in [2]^{n^a} : M'(x, y_1) = 1$. And so $L \in NP = P$, again using the hypothesis.

QED

Exercise 6.2. Prove the following strengthening of Theorem 6.1:

$$\bigcup_d NTime(dn) \subseteq Time(n^{1+\epsilon}) \Rightarrow \bigcup_d \Sigma_i Time(dn) \subseteq Time(n^{1+\epsilon^i}).$$

Exercise 6.3. Show that if $\Sigma_i P = \Pi_i P$ for some i then the PH collapses to $\Sigma_i P$, that is, $PH = \Sigma_i P$.

Theorem 6.2. $NP \subseteq PCKT \Rightarrow PH = \Sigma_2 P$.

Proof. We'll show $\Pi_2 P \subseteq \Sigma_2 P$ and then appeal to Exercise 6.3. Let $f \in \Pi_2 P$ and M be a corresponding machine s.t.

$$f(x) = 1 \Leftrightarrow \forall y_1 \in [2]^{n^d} \exists y_2 \in [2]^{n^d} : M(x, y_1, y_2) = 1.$$

We claim the following equivalent expression for the right-hand side above:

$$\forall y_1 \in [2]^{n^d} \exists y_2 \in [2]^{n^d} : M(x, y_1, y_2) = 1 \Leftrightarrow \exists C \forall y_1 \in [2]^{n^d} : M(x, y_1, C(x, y_1)) = 1,$$

where C ranges over circuits of size $|x|^{d'}$ for some d' . If the equivalence is established the result follows, since evaluating a circuit can be done in power time.

To prove the equivalence, first note that the \Leftarrow direction is obvious, by setting $y_2 := C(x, y_1)$. The interesting direction is the \Rightarrow . We claim that under the assumption, there is a circuit that given x, y_1 outputs a string y_2 that makes $M(x, y_1, y_2)$ accept, if there is such a string.

To verify this, consider the problems CktSat and Search-CktSat which are analogous to the 3Sat and Search-3Sat problems but for general circuits rather than 3CNF. CktSat is in NP, and so by assumption has power-size circuits. By the reduction in Theorem 4.8, Search-CktSat has power-size circuits S as well. Hence, the desired circuit C may, on input x and y_1 produce a new circuit W mapping an input y_2 to $M(x, y_1, y_2)$, and run S on W . **QED**

Exercise 6.4. Prove that $\text{PH} \not\subseteq \text{CktGates}(n^k)$, for any $k \in \mathbb{N}$. (Hint: Existentially guess the truth table of a hard function.)

Improve this to $\Sigma_2\text{P} \not\subseteq \text{CktGates}(n^k)$.

Exercise 6.5. Prove $\text{Exp} \subseteq \text{PCkt} \Rightarrow \text{Exp} = \Sigma_2\text{P}$.

6.2 PH vs. AC

As suggested by the word alternation in Definition 6.1, the power hierarchy PH and its subclasses can be seen as alternating circuits. Before presenting this connection it is useful to write problems in PH in a specific format. The next lemma shows that we can restrict the machine M in Definition 6.1 to read only *one* bit of the input x . The price for this is that we are going to have to introduce an extra quantifier, however this new quantifier will only range over $\log t$ bits.

Lemma 6.1. Let $L \in \Sigma_i\text{P}$. Then there exists a RAM M s.t.:

$$x \in L \Leftrightarrow \exists y_1 \in [2]^{t(|x|)} \forall y_2 \in [2]^{t(|x|)} \dots Q_{i-1} y_{i-1} \in [2]^{t(|x|)} \\ Q_i(y_i, z) \in [2]^{2t(|x|)} Q_{i+1} y_{i+1} \in [2]^{\log t(|x|)} : M(x, y_1, y_2, \dots, y_{i+1}) = 1,$$

and M on input $(x, y_1, y_2, \dots, y_{i+1})$ stops within $ct(|x|)$ steps and only reads one bit of x .

Note the first $i - 1$ quantifiers are over t bits and unchanged from Definition 6.1, the next one is over $2t$ bits, written as a pair (y_i, z) , and the last is over $\log t$. The idea is... *you guessed it!* We are going to guess the bits read from the input, and verify that each of them is correct.

Proof. Let M' be the machine corresponding to L in Definition 6.1. We assume that $Q_i = \exists$, i.e., i is odd. The case $Q_i = \forall$ is left for exercise.

We enlarge Q_i to quantify over additional t bits z which correspond to the bits of x read by M' . The desired machine M simulates M' with the following change. At any time step s , if M' reads bit j of x :

- (1) If $s \neq y_{i+1}$ then M does not read x and instead uses bit z_s ,
- (2) If $s = y_{i+1}$ then M reads the corresponding bit of x , and it also checks that this bit equals z_s . If it doesn't M rejects.

By inspection, M reads at most one bit of x .

Correctness is argued as follows. For any y_1, y_2, \dots, y_i , if M' accepts x then there are values z for the bits read from the input x that cause M to accept. Conversely, if M accepts for some z then this z matches the bits read from x by M' , for else M would reject in (2). Hence M' accepts x as well. **QED**

We now show how to simulate computation in PH by small-depth circuits. The size of the circuit is exponential in the time; such a simulation would not be interesting if the circuit is not explicit and the time is more than the input length, since every function on n bits has circuits of size 2^n (Theorem 2.4). By contrast, the simulation will give explicit circuits and also apply to running times less than the input length. The setting of running time less than the input length makes the simulation interesting even for non-explicit circuits, and is soon to play a critical role.

Lemma 6.2. Any function in $\Sigma_d \text{Time}(t) \cup \Pi_d \text{Time}(t)$ has on inputs of length n alternating circuits of depth $d + 1$ and $2^{ct(n)}$ gates. The fan-in of each gate is $\leq 2^{ct(n)}$ and the fan-in of the gates closest to the input is $\leq t(n)$. Moreover, the circuit is explicit in the following sense: Given an index to a gate g of fan-in h and a number $i \leq h$ we can compute the index of child i of g in $\text{Time}(cn)$.

In fact, most of this circuit only depends on t . The only thing that depends on the actual function being computed are the connections between the gates closest to the input and the input.

Proof. We apply Lemma 6.1. Then an \exists (resp. \forall) quantifier on b bits corresponds to an Or (resp. And) gate with fan-in 2^b . A gate g closest to the input correspond to the computation of the RAM for fixed quantified variables. This computation depends on at most one bit of x . If this computation is a constant independent of x , we simply replace this gate with the appropriate constant. Otherwise, if it depends on a bit x_j of x then the computation of the RAM is either x_j or $\neg x_j$. Thus we can connect gate g to either input gate x_j or input gate $\neg x_j$.

An index to a gate g next to the input is just an assignment to the quantified variables y_i . Given such an index and $i \leq t$ we can compute in linear time which input bit it depends on. This is done by simulating the machine until the i -th time it reads an input bit. Note this simulation runs in time ct which is linear in the length of an index to the gate. **QED**

6.3 BPP in PH

It is not known if BPP is contained in NP. However, we can show that BPP is in PH. More precisely, the following two simulations are known. The first optimizes the number of quantifiers, the second the time. This should be contrasted with various *conditional* results suggesting that in fact a quasilinear deterministic simulation (with no quantifiers) is possible.

Theorem 6.3. For every function t we have:

- (1) $\text{BPTime}(t) \subseteq \Sigma_2\text{Time}(t^2 \log^c t)$, and
- (2) $\text{BPTime}(t) \subseteq \Sigma_3\text{Time}(t \log^c t)$.

A good way to think of these results is as follows. Fix a $\text{BPTime}(t)$ machine M and an input $x \in [2]^n$. Now the alternating machine is trying to decide if for most choices of the random bits y we have $M(x, y) = 1$, or if for most choices we have $M(x, y) = 0$. This can be seen as a version of the Majority problem, with two critical features. The first is that it is *succinct* in the sense of section 5.4.1, that is, the alternating machine does not have access to the exponentially-long majority instance, but rather has access to a small circuit $M(x, \cdot)$ s.t. $M(x, y)$ is bit y of the majority instance. The second is that instances have a *gap*. We define this gap-majority problem next. We define the non-succinct version which is of independent interest, and so use the letter n to indicate input length. But recall that for the application later in this section the input is given succinctly, as just remarked, so the gap majority instance will actually be on a number of bits that is exponential in the input length to the alternating machine.

Definition 6.2. $\text{Gap-Maj}_{\alpha, \beta}$ is the problem of deciding if an input $x \in [2]^n$ has weight $|x| \leq \alpha n$ or $|x| \geq \beta n$.

As discussed in section §6.2, it is useful to think of alternating computation as alternating circuits. Indeed, the circuit result that is the starting point of all these simulations is the following somewhat surprising construction of small-depth alternating circuits for Gap-Maj . By contrast, (non-gap) Maj does not have small constant-depth alternating circuits, as we will prove in Chapter ??.

Lemma 6.3. [5] $\text{Gap-Maj}_{1/3, 2/3}(x)$ has alternating circuits of depth 3 and size n^c . Moreover, the gates at distance 1 from the input have fan-in $\leq c \log n$.

Proof. This is a striking application of the probabilistic method. For a fixed pair of inputs (x, y) we say that a distribution C on circuits *gives* $(\leq p, \geq q)$ if $\mathbb{P}_C[C(x) = 1] \leq p$ and $\mathbb{P}_C[C(y) = 1] \geq q$; and we similarly define *gives* with reverse inequalities. Our goal is to have a distribution that gives

$$(\leq 2^{-n}, \geq 1 - 2^{-n}) \tag{6.1}$$

for every pair $(x, y) \in [2]^n \times [2]^n$ where $|x| \leq n/3$ and $|y| \geq 2n/3$. Indeed, if we have that we can apply a union bound over the $< 2^n$ inputs to obtain a fixed circuit that solves Gap-Maj .

We construct the distribution C incrementally. Fix any pair (x, y) as above. Begin with the distribution C_\wedge obtained by picking $2 \log n$ bits uniformly from the input and computing their And. This gives

$$\left((1/3)^{2 \log n}, (2/3)^{2 \log n} \right).$$

Let $p := (1/3)^{2 \log n}$ and note $(2/3)^{2 \log n} = p \cdot n^2$. So we can say that C_\wedge gives

$$\left(\leq p, \geq p \cdot n^2 \right).$$

Now consider the distribution C_\vee obtained by complementing the circuits in C_\wedge . This gives

$$\left(\geq 1 - p, \leq 1 - p \cdot n^2 \right).$$

Next consider the distribution $C_{\wedge\vee}$ obtained by taking the And of $m := p^{-1}/n$ independent samples of C_\vee . This gives

$$\left(\geq (1 - p)^m, \leq (1 - p \cdot n^2)^m \right).$$

To make sense of these quantities we use the basic approximations

$$1 + \alpha \leq e^\alpha \leq 1 + 2\alpha$$

valid for all $\alpha \in [0, 1]$. These imply $(1 - p)^m \geq e^{-pm/2} = e^{-1/(2n)} \geq 0.9$ and $(1 - p \cdot n^2)^m \leq e^{-n}$. Summarizing, this gives

$$\left(\geq 0.9, \leq e^{-n} \right).$$

Next consider the distribution $C_{\vee\wedge}$ obtained by complementing the circuits in $C_{\wedge\vee}$. This gives

$$\left(\leq 0.1, \geq 1 - e^{-n} \right).$$

Finally, consider the distribution $C_{\wedge\vee\wedge}$ obtained by taking the And of n independent samples of $C_{\vee\wedge}$. This gives

$$\left(\leq 0.1^n, \geq (1 - e^{-n})^n \right).$$

To make sense of the rightmost quantity we can use the approximation

$$(1 + \alpha)^r \geq 1 + r\alpha$$

valid for all $\alpha \geq -1$ and $r \geq 1$. Thus this gives

$$\left(\leq 0.1^n, \geq 1 - ne^{-n} \right).$$

We have $ne^{-n} < 2^{-n}$. Thus this distribution in particular gives equation (6.1). The bounds on the number of gates and the fan-in holds by inspection. **QED**

Exercise 6.6. Prove $\text{Gap-Maj}_{1/2-1/\sqrt{\log n}, 1/2+1/\sqrt{\log n}}$ has alternating circuits of depth c and size n^c .

Exercise 6.7. Assume the circuit in Lemma 6.3 is explicit in the sense of Lemma 6.2. Prove Theorem 6.3.

There remains to construct explicit circuits for Gap-Maj. We give a construction which has worse parameters than Lemma 6.3 but is simple and suffices for (1) in Theorem 6.3. The idea is that if the input weight of x is large, then we can find a few *shifts* of the ones in x that cover each of the n bits. But if the weight of x is small we can't. By "shift" by s we mean the string $x_{i\oplus s}$, obtained from x by permuting the indices by xoring them with s . (Other permutations would work just as well.)

Lemma 6.4. Let $r := \log n$. The following circuit solves $\text{GapMaj}_{1/r^2, 1-1/r^2}$ on every $x \in [2]^n$:

$$\bigvee s_1, s_2, \dots, s_r \in [2]^r : \bigwedge i \in [2]^r : \bigvee j \in \{1, 2, \dots, r\} : x_{i\oplus s_j}.$$

Note that

$$\bigwedge i \in [2]^r : \bigvee j \in \{1, 2, \dots, r\} : x_{i\oplus s_j}$$

means that every bit i is covered by some shift s_j of the input x .

Proof. Assume $|x| \leq n/r^2$. Each shift s_i contributes at most n/r^2 ones. Hence all the r shifts contribute $\leq n/r$ ones, and we do not cover every bit i .

Now assume $|x| \geq n(1 - 1/r^2)$. We show the existence of shifts s_i that cover every bit by the probabilistic method. Specifically, for a fixed x we pick the shifts uniformly at random and aim to show that the probability that we do not cover every bit is < 1 . Indeed:

$$\begin{aligned} & \mathbb{P}_{s_1, s_2, \dots, s_r} [\exists i \in [2]^r : \forall j \in \{1, 2, \dots, r\} : x_{i\oplus s_j} = 0] \\ & \leq \sum_{i \in [2]^r} \mathbb{P}_{s_1, s_2, \dots, s_r} [\forall j \in \{1, 2, \dots, r\} : x_{i\oplus s_j} = 0] && \text{(union bound)} \\ & = \sum_{i \in [2]^r} \mathbb{P}_s [x_{i\oplus s} = 0]^r && \text{(independence of the } s_i) \\ & \leq \sum_{i \in [2]^r} (1/r^2)^r && \text{(by assumption on } |x|) \\ & \leq (2/r^2)^r \\ & < 1, \end{aligned}$$

as desired. **QED**

Exercise 6.8. Prove (1) in Theorem 6.3.

Lemma 6.4 is not sufficient for (2) in Theorem 6.3. One can prove (2) by *derandomizing* the shifts in Lemma 6.4. This means generating their r^2 bits using a seed of only $r \log^c r$ bits (instead of the trivial r^2 in Lemma 6.4.). This is done in section 13.1.3.

Exercise 6.9. Prove:

- (1) $P = NP \Rightarrow P = BPP$.
- (2) $\Sigma_2 P \subseteq BPP \Rightarrow PH$ collapses.

6.4 The quantifier calculus

We have extended P with \exists and \forall quantifiers. We have also extended it with randomness to obtain BPP. As alluded to before, we can also think of BPP as a quantifier BP applied to P. The Unique-3Sat problem (Theorem 5.8) also points to a new quantifier, “exists unique.” We now develop a general calculus of quantifiers, and examine fundamental relationships between them. For simplicity, we only consider power-time computation.

Definition 6.3. Let C be a class of functions mapping $[2]^* \rightarrow [2]$. We define $L \in \text{Op} \cdot C$ if there is $L' \in C$ and $d \in \mathbb{N}$ such that

- $\text{Op} = \text{Maj}$

$$x \in L \Leftrightarrow \mathbb{P}_{y \in [2]^{|x|^d}}[(x, y) \in L'] \geq 1/2.$$

- $\text{Op} = \text{BP}$

$$\begin{aligned} x \in L &\Rightarrow \mathbb{P}_{y \in [2]^{|x|^d}}[(x, y) \in L'] \geq 2/3, \\ x \notin L &\Rightarrow \mathbb{P}_{y \in [2]^{|x|^d}}[(x, y) \in L'] \leq 1/3. \end{aligned}$$

- $\text{Op} = \oplus$ (read: parity)

$$x \in L \Leftrightarrow \text{there is an odd number of } y \in [2]^{|x|^d} : (x, y) \in L'.$$

- $\text{Op} = \exists$

$$x \in L \Leftrightarrow \exists y \in [2]^{|x|^d} : (x, y) \in L'.$$

- $\text{Op} = \forall$

$$x \in L \Leftrightarrow \forall y \in [2]^{|x|^d} : (x, y) \in L'.$$

With this notation we have: $\text{NP} = \exists \cdot \text{P}$, $\text{BPP} = \text{BP} \cdot \text{P}$, $\Sigma_2\text{P} = \exists \cdot \forall \cdot \text{P}$.

6.5 PH is a random low-degree polynomial

In this section we prove the following result.

Theorem 6.4. $\text{PH} \subseteq \text{BP} \cdot \oplus \cdot \text{P}$.

This is saying that any constant number of \exists and \forall quantifier can be replaced by a BP quantifier followed by a \oplus quantifier. Let's see what this has to do with the title of this section. Where is the polynomial? Consider polynomials over $\mathbb{F}_2 \in [2]$. Recall that such a polynomial over n bits is an object like

$$p(x_1, x_2, \dots, x_n) = x_1 \cdot x_2 + x_3 + x_7 \cdot x_2 \cdot x_1 + x_2 + 1.$$

Because we are only interested in inputs in $[2]$ we have $x^i = x$ for any $i \geq 1$ and any variable x , so we don't need to raise variables to powers bigger than one.

Example 6.1. The And function on n bits can be written as the polynomial

$$\text{And}(x_1, x_2, \dots, x_n) = \prod_{i=1}^n x_i.$$

The Or function on n bits can be written as the polynomial

$$\text{Or}(x_1, x_2, \dots, x_n) = 1 + \text{And}(1 + x_1, 1 + x_2, \dots, 1 + x_n) = 1 + \prod_{i=1}^n (1 + x_i).$$

For $n = 2$ we have

$$\text{Or}(x_1, x_2) = x_1 + x_2 + x_1 \cdot x_2.$$

The polynomial corresponding to a PH computation will have an exponential number of terms, so we can't write it down. The big sum over all its monomials corresponds to the \oplus in Theorem 6.4. The polynomial will be sufficiently explicit: we will be able to compute each of its monomials in P. Finally, there won't be just one polynomial, but we will have a distribution on polynomials, and that's the BP part.

Confusing? Like before, a good way to look at this result is in terms of alternating circuits. We state the basic circuit result behind Theorem 6.4 after a definition. The result is of independent interest and will be useful later in Chapter 10.

Definition 6.4. A distribution P on polynomials computes a function $f : [2]^n \rightarrow [2]$ with error ϵ if for every x we have

$$\mathbb{P}_P[P(x) = f(x)] \geq 1 - \epsilon.$$

Theorem 6.5. Let $C : [2]^n \rightarrow [2]$ be an alternating circuit of depth d and size s . Then there is a distribution P on polynomials over \mathbb{F}_2 of degree $\log^{d-1} s/\epsilon$ that computes C with error ϵ .

Ultimately we only need constant error, but the construction requires small error. Jumping ahead, this is because we construct distributions for each gate separately, and we need the error to be small enough for a union bound over all gates in the circuit.

The important point in Theorem 6.5 is that if the depth d is small (e.g., constant) (and the size is not enormous and the error is not too small) then the degree is small as well. For example, for power-size alternating circuits of constant depth the degree is power logarithmic for constant error.

Let us slowly illustrate the ideas behind Theorem 6.5 starting with the simplest case: C is just a single Or gate on n bits.

Lemma 6.5. For every ϵ and n there is a distribution P on polynomials of degree $\log 1/\epsilon$ in n variables over \mathbb{F}_2 that computes Or with error ϵ .

Proof. For starters, pick the following distribution on linear polynomials: For a uniform $A = (A_1, A_2, \dots, A_n) \in [2]^n$ output the polynomial

$$p_A(x_1, x_2, \dots, x_n) := \sum_i A_i \cdot x_i.$$

Let us analyze how p_A behaves on a fixed input $x \in [2]^n$:

- If $\text{Or}(x) = 0$ then $p_A(x) = 0$;
- If $\text{Or}(x) = 1$ then $\mathbb{P}_A[p_A(x) = 1] \geq 1/2$.

While the error is large in some cases, a useful feature of p_A is that it never makes mistakes if $\text{Or}(x) = 0$. This allows us to easily reduce the error by taking $t := \log 1/\epsilon$ polynomials p_A and combining them with an Or.

$$p_{A_1, A_2, \dots, A_t}(x) := p_{A_1}(x) \vee p_{A_2}(x) \vee \dots \vee p_{A_t}(x).$$

The analysis is like before:

- If $\text{Or}(x) = 0$ then $p_{A_1, A_2, \dots, A_t}(x) = 0$;
- If $\text{Or}(x) = 1$ then $\mathbb{P}_{A_1, A_2, \dots, A_t}[p_{A_1, A_2, \dots, A_t}(x) = 1] \geq 1 - (1/2)^t \geq 1 - \epsilon$.

It remains to bound the degree. Each p_{A_i} has degree 1. The Or on t bits has degree t by Example 6.1. Hence the final degree is $t = \log 1/\epsilon$. **QED**

Exercise 6.10. Obtain the same result for $C = \text{And}$.

Now we would like to handle general circuits which have any number of And and Or gates. As mentioned earlier, we apply the construction above to every gate, and compose the polynomials. We pick the error at each gate small enough so that we can do a union bound over all gates.

Proof of Theorem 6.5. We apply Lemma 6.5 to every gate in the circuit with error ϵ/s . By a union bound, the probability that any gate makes a mistake is ϵ , as desired.

The final polynomial is obtained by composing the polynomials of each gate. The composition of a polynomial of degree d_1 with another of degree d_2 results in a polynomial of degree $d_1 \cdot d_2$. Since each polynomial has degree $\log s/\epsilon$, and we compose $d - 1$ times, the final degree is $\log^{d-1} s/\epsilon$. **QED**

6.5.1 Back to PH

We have proved Theorem 6.5 which is a circuit analogue of Theorem 6.4. We now go back to the PH. First, we have to come up with a more explicit description of the polynomials, instead of picking them at random. This is similar to the way we proceeded in section §6.3:

After a non-explicit construction (Lemma 6.3) we then obtained an explicit construction (Lemma 6.4).

Let us go back to the simplest case of Or. Recall that the basic building block in the proof of Lemma 6.5 was the construction of a distribution p_A on linear polynomials which are zero on the all-zero input (which just means that they do not have constant terms), and are often non-zero on any non-zero input. We introduce a definition, since now we will have several constructions with different parameters.

Definition 6.5. A distribution p_A on linear polynomials with no constant term has the Or property if $\mathbb{P}_A[p_A(x) = 1] \geq 1/3$ for any $x \neq 0$. We identify p_A with the n bits A corresponding to its coefficients.

The next lemma shows that we can compute distributions on linear polynomials with the Or property from a seed of just $\log n + c$ bits, as opposed to the n bits that were used for A in the proof of Lemma 6.5. This important fact is generalized and put in context in section 13.1.2. Recall that for our application to Lemma 6.5 the polynomials have an exponential number of monomials and so we cannot afford to write them down. Instead we shall guarantee that given a seed r and an index to a monomial we can compute the monomial via a function f in P. In this linear case, for a polynomial in n variables we have $\leq n$ monomials x_i . So the function f takes as input r and a number $i \leq n$ and outputs the coefficient to x_i .

Lemma 6.6. Given n , $i \leq n$, and $r \in [2]^{2^{\log n+c}}$ we can compute in P a function $f(r, i)$ such that for uniform $R \in [2]^{2^{\log n+c}}$ the distribution

$$(f(R, 1), f(R, 2), \dots, f(R, n))$$

has the Or property.

Proof. Let $q := 2^{\log n+c}$ and identify the field \mathbb{F}_q with bit strings of length $\log q$. We view r as a pair $(s, t) \in (\mathbb{F}_q)^2$. Then we define

$$f((s, t), i) := \langle s^i, t \rangle$$

where s^i is exponentiation in \mathbb{F}_q and $\langle \cdot, \cdot \rangle : (\mathbb{F}_q)^2 \rightarrow [2]$ is defined as $\langle u, v \rangle := \sum u_i \cdot v_i$ over \mathbb{F}_2 .

To show that this has the Or property, pick any non-zero $x \in [2]^n$. We have to show that

$$p := \mathbb{P}_{S,T}[\sum_i \langle S^i, T \rangle x_i = 1] \geq 1/3.$$

The critical step is to note that

$$\sum_i \langle S^i, T \rangle x_i = \sum_i \langle x_i \cdot S^i, T \rangle = \langle \sum_i x_i \cdot S^i, T \rangle.$$

Now, if $x \neq 0$, then the probability over S that $\sum_i x_i \cdot S^i = 0$ is $\leq n/q \leq 1/6$. This is because any S that gives a zero is a root of the non-zero, univariate polynomial $q(y) := \sum_i x_i \cdot y^i$ of degree $\leq n$ over \mathbb{F}_q , and so the bound follows by Lemma 2.3.

Whenever $\sum_i x_i \cdot S^i \neq 0$, the probability over T that $\langle \sum_i x_i \cdot S^i, T \rangle = 0$ is $1/2$. Hence our target probability p above satisfies

$$p \geq 1/2 - 1/6$$

as desired. **QED**

Exercise 6.11. Give an alternative construction of a distribution with the Or property following this guideline.

(1) Satisfy the Or property for every input x with weight 1.

(2) For any j , satisfy the Or property for every input x with weight between 2^j and 2^{j+1} .

Use Lemma 5.3.

(3) Combine (2) with various j to satisfy the Or property for every input.

(4) State the seed length for your distribution and compare it to that of Lemma 6.6.

With this in hand, we can now reduce the error in the same way we reduced it in the proof of Lemma 6.5.

Lemma 6.7. Given n , a seed $r \in [2]^{c \log(1/\epsilon) \log n}$, and $m \leq n^{c \log 1/\epsilon}$ we can compute in P a monomial $X_{r,m}$ of degree $c \log 1/\epsilon$ such that the distribution

$$\sum_m X_{R,m}$$

for uniform R computes Or with error ϵ .

Proof. We use the construction

$$p_{A_1, A_2, \dots, A_t}(x) := p_{A_1}(x) \vee p_{A_2}(x) \vee \dots \vee p_{A_t}(x)$$

from the proof of Lemma 6.5, except that each A_i is now generated via Lemma 6.5, and that $t = c \log 1/\epsilon$ (as opposed to $t = \log 1/\epsilon$ before). The bound on the degree is the same as before, as is the proof that it computes Or: The error will be $(1/3)^{c \log 1/\epsilon} \leq \epsilon$.

There remains to show that the monomials can be computed in P. For this we can go back to the polynomial for Or in Example 6.1. Plugging that gives

$$p_{A_1, A_2, \dots, A_t}(x) = \sum_{a \in [2]^t: a \neq 0} \prod_{i \leq t} (p_{A_i}(x) + a_i + 1).$$

We can use m to index a choice of a and then a choice for a monomial in each of the t linear factors $p_{A_i}(x) + a_i + 1$. For each factor we can use Lemma 6.6 to compute the monomials. **QED**

We can now use Lemma 6.7 to prove Theorem 6.4. As in the proof Theorem 6.5 we will convert each gate to a polynomial.

Proof of Theorem 6.4. Let $L \in \text{PH}$. Apply Lemma 6.2 to obtain a corresponding alternating circuit C of depth d and size $s \leq 2^{n^d}$ for a constant d . Replace each gate with the distribution on polynomials given by Lemma 6.7. Note each of these polynomials is on $2^{n^{c_d}}$ variables and has degree n^{c_d} . We set the error ϵ in Lemma 6.7 to be $1/(3s)$. This guarantees that the seed length in each application of Lemma 6.7 is $\leq n^{c_d}$. Moreover, the seed used to sample the polynomials is re-used across all gates. We can afford this because we use a union bound in the analysis. Hence the seed length for all the polynomials is again just n^{c_d} . We can quantify over this many bits using the BP quantifier.

Finally, we have to compose the polynomials at each gate. We are going to show how we can compute the monomials of the composed polynomial in the same way as we computed monomials in Lemma 6.7. Once we have a monomial in the input variables x_1, x_2, \dots, x_n we simply evaluate it in \mathbb{P} on the input bits.

Start at the output gate. We use n^{c_d} bits in the \oplus quantifier to choose a monomial in the corresponding polynomial. We write down this monomial, using Lemma 6.7. This monomial is over the $2^{n^{c_d}}$ variables z_1, z_2, \dots corresponding to the children of the output gate, and only contains n^{c_d} variables. To each z_i there corresponds a polynomial p_i . Choose a monomial from p_i and replace z_i with that monomial; do this for every i . The choice of the monomials is done again using bits quantified by the \oplus quantifier. We continue in this way until we have monomials just in the input variables x_1, x_2, \dots, x_n . Because each monomial is over n^{c_d} variables, and the depth is constant, the total number of bits in the \oplus quantifier, which are needed to choose monomials, is n^{c_d} . **QED**

Exercise 6.12. A set $C \subseteq [2]^n$ of size 2^k is called *linear* if there exists an $n \times k$ matrix M over \mathbb{F}_2 such that $C = \{Mx : x \in [2]^k\}$.

1. Recall error-correcting codes from Exercise 2.21. Prove that a linear set C is an error-correcting code iff the weight of any non-zero string in C is at least $n/3$.
2. Prove the existence of linear error-correcting codes matching the parameters in Exercise 2.21.
3. Let S be a subset of $[2]^k$ s.t. the uniform distribution over S has the Or property. Define $|S| \times k$ matrix M_S where the rows are the elements of S . Prove that $\{M_S x : x \in [2]^k\}$ is an error-correcting code.
4. Give explicit error-correcting codes over ck^2 bits of size 2^k .
5. This motivates improving the parameters of distributions with the Or property. Improve the seed length in Lemma 6.6 to $\log n + c \log \log n$. Hint: What property you need from T ?
6. Give explicit error-correcting codes over $k \log^c k$ bits of size 2^k .

6.6 The power of majority

Exercise 6.13. [The power of majority] Prove:

1. $\text{BPP} \cdot \text{P} \subseteq \text{Maj} \cdot \text{P}$.
2. $\{(M, x, t) : \text{the number of } y \in [2]^{|t|} \text{ such that } M(x, y) = 1 \text{ is } \geq t\} \in \text{Maj} \cdot \text{P}$.
3. The same as 2. but with \geq replaced by \leq .
4. $\text{NP} \subseteq \text{Maj} \cdot \text{P}$.
5. $\text{Maj} \cdot \oplus \cdot \text{P} \subseteq \text{Maj} \cdot \text{Maj} \cdot \text{P}$. (Hint: This is confusing if you don't have the right angle, otherwise is not too bad. For starters, forget everything and imagine you have an integer w in some range and you want to know if w is odd by just asking questions of the type $w \geq t$ and $w \leq t'$, for various t, t' . You want that the number of questions with answer "yes" only depends on whether w is odd or even.)
6. $\text{PH} \subseteq \text{Maj} \cdot \text{Maj} \cdot \text{P}$.

It is not known if NP has linear-size circuits. We saw in Exercise 6.4 that PH does not have circuits of size n^k , for any k . By Exercise 6.13 this holds for $\text{Maj} \cdot \text{Maj} \cdot \text{P}$ as well. The following result improves this $\text{Maj} \cdot \text{P}$. It is particularly interesting because it cannot be established using a well-studied class of techniques which includes all the results about PH we have encountered so far, specifically black-box techniques discussed in section §16.1.

Theorem 6.6. $\text{Maj} \cdot \text{P} \not\subseteq \text{CktGates}(n^k)$, for any $k \in \mathbb{N}$.

6.7 Problems

Problem 6.1. Prove $\Sigma_i \text{Time}(n^\alpha) \not\subseteq \Pi_i \text{Time}(n^\beta)$, for any constants i and any $\alpha > \beta$.

Feel free to work with partial functions and assume $\beta > 1$ if you prefer, because these details are not the main point of the problem.

Problem 6.2. Prove that $\text{QLin-NTime} = \text{QLin-Time} \Rightarrow \text{QLin-BPTime} = \text{QLin-Time}$. (Cf. Definition 5.4, the definition of QLin-BPTime is analogous.)

Problem 6.3. [Arithmetic on truth tables] Let $f : [2]^* \rightarrow [2]$ be a function. Denote by $a_f(n)$ the 2^n -bit integer whose binary representation are the evaluations of f on every input of n bits. For example,

$$a_f(2) = f(11)f(10)f(01)f(00).$$

Define f' to be the function s.t. $a_{f'}(n) = a_f(n) + 1 \pmod{2^{2^n}}$, for every n .

Show that if $f \in \text{PH}$ then $f' \in \text{PH}$.

Problem 6.4. Prove Theorem 6.5 with the bound on the degree replaced by $(\log^{d-1} cs)c \log 1/\epsilon$ (which is better when ϵ is small).

Problem 6.5. (1) Show that for any constants $\alpha < \beta$, $\text{Gap-Maj}_{\alpha,\beta}$ has alternating circuits of depth $c_{\alpha,\beta}$ and size $n^{c_{\alpha,\beta}}$.

(2) Show that $\text{Gap-Maj}_{1/2-1/\log n, 1/2+1/\log n}$ has alternating circuits of depth c and size n^c . (This can be generalized by replacing the \log with \log^a and the c with c_a , but you are not asked to prove this.)

Problem 6.6. Prove $\exists \cdot \text{BP} \cdot \text{P} \subseteq \text{Maj} \cdot \text{P}$.

6.8 Notes

The PH was identified in [182], where Theorem 6.1 is also proved. Theorem 6.2 is from [110].

The connection between PH and AC was pointed out in [61].

The first item in the simulation of BPP Theorem 6.3 is from [177]; the second is from [206].

Maj P was defined in [96].

Parity P was defined in [150].

Theorem 6.4 is from [191].

Theorem 6.5 is from [158].

Lemma 6.6 is from [138], where they actually prove stronger results, see Theorem 13.3. The proof we presented is a particularly nice one from the three presented in the follow-up [13].

Theorem 6.6 is from [202].

For a compendium of problems complete for various levels of the PH, in the style of [64], see [168].

Chapter 7

Space

As mentioned in Chapter 2, Time is only one of several resources we wish to study. Another important one is *space*, which is another name for the *memory* of the machine. Computing under space constraints may be less familiar to us than computing under time constraints, and many surprises lay ahead in this chapter that challenge our intuition of efficient computation.

If only space is under consideration, and one is OK with a constant-factor slackness, then TMs and RAMs are equivalent; much like P is invariant under power changes in time. In a sense, changing the space by a constant factor is like changing the time by a power; from this point of view the equivalence is not surprising.

We shall consider both space bounds bigger than the input length and smaller. For the latter, we have to consider the input separately. The machine should be able to *read* the input, but not *write* on its cells. One way to formalize this is to consider 2TMs, where one tape holds the input and is *read-only*. The other is a standard *read-write tape*.

We also want to compute functions f whose output is more than 1 bit. One option is to equip the machine with yet another tape, which is *write-only*. We prefer to stick to two tapes and instead require that given x, i the i output bit of $f(x)$ is computable efficiently.

Definition 7.1. A function $f : X \subseteq [2]^* \rightarrow [2]^*$ is computable in $\text{Space}(s(n))$ if there is a 2TM which on input (x, i) on the first tape, where $x \in X$ and $i \leq |f(x)|$, outputs the i bit of $f(x)$, and the machine never writes on the first tape and never uses more than $s(|x|)$ cells on the second.

We define:

$$\begin{aligned} L &:= \bigcup_d \text{Space}(d \log n), \\ \text{PSpace} &:= \bigcup_d \text{Space}(n^d). \end{aligned}$$

We investigate next some basic relationship between space and time.

Theorem 7.1. For every functions t and s :

1. $k\text{TM-Time}(t) \subseteq \text{Space}(c_k t)$,
2. $\text{Time}(t) \subseteq \text{Space}(ct \log(nt))$,
3. $\text{Space}(s) \subseteq 2\text{TM-Time}(c^{s(n)} \cdot n^c)$.

Proof. 1. A TM running in time t can only move its heads at most t tape cells. We can write all these contents in one tape. To simulate one step of the k TM we do a pass on all the contents and update them accordingly.

2. A RAM running in time t can only access $\leq t$ memory cells, each containing at most $c \log nt$ bits; the factor n is to take into account that the machine starts with word length $\geq \log n$. We simulate this machine and for each Write operation we add a record on the tape with the memory cell index and the content, similar to Theorem 2.7. When the RAM reads from memory, we scan the records and read off the memory values from one of them. If the record isn't found, then the simulation returns the value corresponding to the initial configuration.

We also allocate $c \log nt$ bits for the registers of the RAM. It can be shown that the operations among them can be performed in the desired space, since they only involve a logarithmic number of bits. A stronger result is proved later in Theorem 7.5.

3. On an input x with $|x| = n$, a $\text{Space}(s)$ machine can be in at most $n^c \cdot c^{s(n)}$ configurations. The first n^c factor is for the head position on the input. The second factor is for the contents of the second tape. Since the machine ultimately stops, configurations cannot repeat, hence the same machine (with no modification) will stop in the desired time. **QED**

A non-trivial saving is given by the following theorem.

Theorem 7.2. $k\text{-TM-Time}(t) \subseteq \text{Space}(c_k t / \log t)$, for every function t .

This result is not specific to MTMs but can be proved for other models such as RAMs.

Proof sketch of Theorem 7.2. We only give a sketch of the beautiful proof here. For simplicity, let us think of 1TMs, for which the problem is already non-trivial. The solution will generalize rather straightforwardly to more tapes.

Let $b := t^c$. Divide the tape into consecutive blocks of $b = t^c$ symbols, and assume that the machine so that it only crosses a block at times that are multiples of b . One can modify machines to have this property by increasing the number of tapes and only paying a constant-factor overhead in time.

We also divide time into t/b blocks. For $i \leq t/b$ we talk of simulating the machine up to time block i . This means computing state and head positions right at the end of time block i , and also the contents of the block the machine worked on in time block i .

We construct a graph with $c_M \cdot t \cdot t/b$ nodes. Here a node encodes a time block i , the state of the machine right at the end of that time block, and the position of the head on the tape.

We always have edges $i - 1 \rightarrow i$. In addition, we place an edge $i \rightarrow j$ if the machine in time block j is working on the same block the machine was working in time block i , and i is the largest index less than j with this property.

The main beautiful idea is to *pebble* this graph. The rules for such a pebbling are that we can place a pebble on any node whose predecessors are all pebbled, and remove a pebble from any node. It is known that if t is reachable from s in a graph with n nodes then we can pebble t using $\leq cn/\log n$ pebbles.

If you can pebble i , then you can simulate the machine up to time block i , using space that is roughly the number of pebbles times the size of a block. The edges $i - 1 \rightarrow i$ will be needed to know the state and head positions, and the edges $i \rightarrow j$ will be needed in case the machine is crossing a block.

There remains to compute the pebbling. This is a non-trivial task, since the pebbling can involve an exponential number of moves, but we can use a recursive strategy to solve the following problem (similar to Theorem 7.17): Given a current pebbling of the graph, compute the next move in an optimal pebbling. This can be done in space roughly square the size of the pebbling, which will be within our budget by our choice of b . **QED**

Exercise 7.1. Improve Theorem 7.2 for 1TMs.

From Theorem 7.1 and the next exercise we have

$$L \subseteq P \subseteq NP \subseteq PH \subseteq PSpace \subseteq Exp.$$

Exercise 7.2. Prove $PH \subseteq PSpace$.

Just like for Time, for space one has universal machines and a hierarchy theorem. The latter implies $L \neq PSpace$. Hence, analogously to the situation for Time and NTime (section §5.1), we know that at least one of the inclusions above between L and $PSpace$ is strict. Most people seem to think that all are, but nobody can prove that any specific one is.

7.1 Branching programs

Branching programs are the non-uniform counterpart of Space, just like circuits are the non-uniform counterpart of Time.

Definition 7.2. A (*branching*) *program* is a directed acyclic graph. A node can be labeled with an input variable, in which case it has two outgoing edges labeled 0 and 1. Alternatively a node can be labeled with 0 or 1, in which case it has no outgoing edges. One special node is the *start* node.

The *space* of the program with S nodes is $\log S$. A program computes a function $f : [2]^n \rightarrow [2]$ by following the path from the starting node, following edge labels corresponding to the input, and outputting $b \in [2]$ as soon as it reaches a node labeled b .

Theorem 7.3. Suppose a 2TM M computes $f : [2]^n \rightarrow [2]$ in space s . Then f has branching programs of space $c_M(s + \log n)$. In particular, any $f \in L$ has branching programs of size n^{c_f} .

Proof. Each node in the program corresponds to a configuration. **QED**

Definition 7.3. The branching program given by Theorem 7.3 is called the *configuration graph* of M .

The strongest available impossibility result for branching programs is the following.

Theorem 7.4. There are explicit functions (in P) that require branching programs of size $cn^2/\log n$ on inputs of length n .

7.2 The power of L

Computing with severe space bounds, as in L, seems quite difficult. Also, it might be somewhat less familiar than, say, computing within a time bound. It turns out that L is a powerful class capable of amazing computational feats that challenge our intuition of efficient computation. Moreover, these computational feats hinge on deep mathematical techniques of wide applicability. We hinted at this in Chapter 1. We now give further examples. At the same time we develop our intuition of what is computable with little space.

To set the stage, we begin with a composition result. In the previous sections we used several times the simple result that the composition of two maps in P is also in P. This is useful as it allows us to break a complicated algorithm in small steps to be analyzed separately – which is a version of the *divide et impera* paradigm. A similar composition result holds and is useful for space, but the argument is somewhat less obvious.

Lemma 7.1. Let $f_1 : [2]^* \rightarrow [2]^*$ be in $\text{Space}(s_1)$ and satisfy $|f_1(x)| \leq m(|x|)$ for a function m . Suppose $f_2 : [2]^* \rightarrow [2]$ is in $\text{Space}(s_2)$.

Then the composed function g defined as $g(x) = f_2(f_1(x))$ is computable in space $c(s_2(m(n)) + s_1(n) + \log nm(n))$.

In particular, if f_1 and f_2 are in L then g is in L, as long as $m \leq n^d$ for a constant d .

Exercise 7.3. Prove this.

7.2.1 Arithmetic

A first example of the power of L is given by its ability to perform basic arithmetic. Grade school algorithms use a lot of space, for example they employ space $\geq n$ to multiply two n -bit integers.

Theorem 7.5. The following problems are in L:

1. Addition of two input integers.
2. Iterated addition: Addition of any number of input integers.
3. Multiplication of two input integers.

4. Iterated multiplication: Multiplication of any number of input integers.
5. Division of two integers.

Iterated multiplication is of particular interest because it can be used to compute “pseudorandom functions.” Such objects shed light on our ability to prove impossibility results via the “Natural Proofs” connection which we will see in Chapter 19.

Proof of 1. in Theorem 7.5. We are given as input $x, y \in \mathbb{N}$ and an index i and need to compute bit i of $x + y$. Starting from the least significant bits, we add the bits of x and y , storing the carry of 1 bit in memory. Output bits are discarded until we reach bit i , which is output. **QED**

Exercise 7.4. Prove 2. and 3. in Theorem 7.5.

Proving 4. and 5. is more involved and requires some of those deep mathematical tools of wide applicability we alluded to before. Division can be computed once we can compute iterated multiplication. In a nutshell, the idea is to use the expansion

$$\frac{1}{x} = \sum_{i \geq 0} (-1)^i (x - 1)^i.$$

We omit details about bounding the error. Instead, we point out that this requires computing powers $(x - 1)^i$ which is an example of iterated multiplication (and in fact is no easier).

So for the rest of this section we focus on iterated multiplication. Our main tool for this is the Chinese-remaindering representation of integers, abbreviated CRR.

Definition 7.4. We denote by \mathbb{Z}_m the integers modulo m equipped with addition and multiplication (modulo m).

Theorem 7.6. Let p_1, \dots, p_ℓ be distinct primes and $m := \prod_i p_i$. Then \mathbb{Z}_m is isomorphic to $\mathbb{Z}_{p_1} \times \dots \times \mathbb{Z}_{p_\ell}$.

The forward direction of the isomorphism is given by the map

$$x \in \mathbb{Z}_m \rightarrow (x \bmod p_1, x \bmod p_2, \dots, x \bmod p_\ell) \in \mathbb{Z}_{p_1} \times \dots \times \mathbb{Z}_{p_\ell}.$$

For the converse direction, there exist integers $e_1, \dots, e_\ell \leq (p')^c$, depending only on the p_i such that the converse direction is given by the map

$$(x \bmod p_1, x \bmod p_2, \dots, x \bmod p_\ell) \in \mathbb{Z}_{p_1} \times \dots \times \mathbb{Z}_{p_\ell} \rightarrow x := \sum_{i=1}^{\ell} e_i \cdot (x \bmod p_i).$$

Each integer e_i is 0 mod p_j for $j \neq i$ and is 1 mod p_i .

Example 7.1. \mathbb{Z}_6 is isomorphic to $\mathbb{Z}_2 \times \mathbb{Z}_3$. The equation $2 + 3 = 5$ corresponds to $(0, 2) + (1, 0) = (1, 2)$. The equation $2 \cdot 3 = 6$ corresponds to $(0, 2) + (1, 0) = (0, 0)$. Note how addition and multiplication in CRR are performed in each coordinate separately; how convenient.

To compute iterated multiplication the idea is to move to CRR, perform the multiplications there, and then move back to standard representation. A critical point is that each coordinate in the CRR has a representation of only $c \log n$ bits, which makes it easy to perform iterated multiplication one multiplication at the time, since we can afford to write down intermediate products.

The algorithm is as follows:

Computing the product of input integers x_1, \dots, x_t .

1. Let $\ell := n^c$ and compute the first ℓ prime numbers p_1, p_2, \dots, p_ℓ .
2. Convert the input into CRR: Compute $(x_1 \bmod p_1, \dots, x_1 \bmod p_\ell), \dots, (x_t \bmod p_1, \dots, x_t \bmod p_\ell)$.
3. Compute the multiplications in CRR: $(\prod_{i=1}^t x_i \bmod p_1), \dots, (\prod_{i=1}^t x_i \bmod p_\ell)$.
4. Convert back to standard representation.

Exercise 7.5. Prove the correctness of this algorithm.

Now we explain how steps 1, 2, and 3 can be implemented in L. Step 4 can be implemented in L too, but showing this is somewhat technical due to the computation of the numbers e_i in Theorem 7.6. However these numbers only depend on the input length, and so we will be able to give a self-contained proof that iterated multiplication has branching programs of size n^c .

Step 1

By Theorem 2.14, the primes p_i have magnitude $\leq n^c$ and so can be represented with $c \log n$ bits. We can enumerate over integers with $\leq c \log n$ bits in L. For each integer x we can test if it's prime by again enumerating over all integers y and z with $\leq c \log n$ bits and checking if $x = yz$, say using the space-efficient algorithm for multiplication in Theorem 7.5. (The space required for this step would in fact be $c \log \log n$.)

Step 2

We explain how given $y \in [2]^n$ we can compute $(y \bmod p_1, \dots, y \bmod p_\ell)$ in L. If y_j is bit j of y we have that

$$\begin{aligned} y \bmod p_i &= \left[\sum_{j=0}^{n-1} (2^j y_j) \right] \bmod p_i \\ &= \left[\sum_{j=0}^{n-1} (2^j \bmod p_i) y_j \right] \bmod p_i. \end{aligned}$$

Note that the values $a_{i,j} := 2^j \bmod p_i$ can be computed in L and only take $c \log n$ bits. Multiplying by y_j is also in L. Hence the problem reduces to iterated addition of n numbers which is in L by Theorem 7.5.

Step 3

This is a smaller version of the original problem: for each $j \leq \ell$, we want to compute $(\prod_{i=1}^t x_i \bmod p_j)$ from $x_1 \bmod p_j, \dots, x_t \bmod p_j$. However, as mentioned earlier, each $(x_i \bmod p_j)$ is at most n^c in magnitude and thus has a representation of $c \log n$ bits. Hence we can just perform one multiplication at the time in L.

Step 4

By Theorem 7.6, to convert back to standard representation from CRR we have to compute the map

$$(y \bmod p_1, \dots, y \bmod p_\ell) \rightarrow \sum_{i=1}^{\ell} e_i \cdot (y \bmod p_i).$$

Assuming we can compute the e_i , this is just multiplication and iterated addition, which are in L by Theorem 7.5.

Putting the steps together

Combining the steps together we can compute iterated multiplication in L as long as we are given the integers e_i in Theorem 7.6.

Theorem 7.7. Given integers x_1, x_2, \dots, x_t , and given the integers e_1, e_2, \dots, e_ℓ as in Theorem 7.6, where $\ell = n^c$, we can compute $\prod_i x_i$ in L.

In particular, because the e_i only depend on the input length, but not on the x_i they can be hardwired in a branching program.

Corollary 7.1. Iterated multiplication has branching programs of size n^c .

Exercise 7.6. Show that given integers x_1, x_2, \dots, x_t and y_1, y_2, \dots, y_t one can decide if

$$\prod_{i=1}^t x_i = \prod_{i=1}^t y_i$$

in L. You cannot use the fact that iterated multiplication is in L, a result which we stated but not completely proved.

Exercise 7.7. Show that the iterated multiplication of $d \times d$ matrices over the integers has branching programs of size n^{cd} .

7.2.2 Graphs

We now give another example of the power of L.

Definition 7.5. The undirected reachability problem: Given an undirected graph G and two nodes s and t in G determine if there is a path from s to t .

Standard time-efficient algorithms to solve this problem *mark* nodes in the graph. In logarithmic space we can keep track of a constant number of nodes, but it is not clear how we can avoid repeating old steps forever.

Theorem 7.8. Undirected reachability is in L.

The idea behind this result is that a *random walk* on the graph will visit every node, and can be computed using small space, since we just need to keep track of the current node. Then, one can *derandomize* the random walk and obtain a deterministic walk, again computable in L.

7.2.3 Linear algebra

Our final example comes from linear algebra. Familiar methods for solving a linear system

$$Ax = b$$

can be done requires a lot of space. For example using elimination we need to rewrite the matrix A . Similarly, we cannot easily compute determinants using small space. However, a different method exists.

Theorem 7.9. Solving a linear system is computable in $\text{Space}(c \log^2 n)$.

7.3 Checkpoints

The *checkpoint* technique is a fundamental tool in the study of space-bounded computation. Let us illustrate it at a high level. Let us consider a graph G , and let us write $u \rightsquigarrow^t v$ if there is a path of length $\leq t$ from u to v . The technique allows us to *trade the length of the path with quantifiers*. Specifically, for any parameter b , we can break down paths from u to v in b smaller paths that go through $b - 1$ checkpoints. The length of the smaller paths needs be only t/b (assuming that b divides t). We can guess the breakpoints and verify each smaller path separately, at the price of introducing quantifiers but with the gain that the path length got reduced from t to t/b . The checkpoint technique is thus an instantiation of the general paradigm of guessing computation and verifying it locally, introduced in Chapter 5. One difference is that now we are only going to guess *parts* of the computation.

The checkpoint technique

$$u \rightsquigarrow^t v \Leftrightarrow \exists p_1, p_2, \dots, p_{b-1} : \forall i \in \{0, 1, b-1\} : p_i \rightsquigarrow^{t/b} p_{i+1},$$

where we denote $p_0 := u$ and $p_b := v$.

An important aspect of this technique is that it can be applied *recursively*: We can apply it again to the problems $p_i \rightsquigarrow^{t/b} p_{i+1}$. We need to introduce more quantifiers, but we can reduce the path length to t/b^2 , and so on. We will see several instantiations of this technique, for various settings of parameters, ranging from $b = 2$ to $b = n^c$.

We now utilize the checkpoint technique to show a simulation of small-space computation by small-depth alternating circuits.

Theorem 7.10. A function computable by a branching program with S nodes is also computable by an alternating circuit of depth $c \log_b S$ and size $S^{b \log_b S + c}$, for any $b \leq S$.

To illustrate the parameters, suppose $S = n^a$, and let us pick $b := n^\epsilon$ where $n \geq c_{a,\epsilon}$. Then we have AltCkts of depth $d := ca/\epsilon$ and size $\leq S^{n^\epsilon d + c} = 2^{n^{c\epsilon}}$. In other words, we can have depth d and size $2^{n^{ca/d}}$, for every d . Another way of saying this is that the circuit has constant depth and power size on inputs of power-logarithmic length. This is quite useful to design AltCkts

Exercise 7.8. Prove that for any t and d the Majority function on $\log^d t$ bits can be computed by AltCkts of size t^{c_d} and depth c_d .

The parameters in the above discussion before the exercise in particular hold for every function in L. We will later give explicit functions (also in P) which cannot be computed by AltCkts of depth d and size $2^{n^{c/d}}$, “just short” of ruling out L. This state of affairs is worth emphasis:

- (1) Every f in L has alternating circuits of depth d and size $2^{n^{c_f/d}}$.
- (2) We can prove that there are explicit functions (also in L) which cannot be computed by circuits of depth d and size $2^{n^{c/d}}$.
- (3) Improving the constant in the double exponent for a function in P would yield $L \neq P$. In this sense, the result in (2) is the best possible short of proving a major separation.

Proof. We apply the checkpoint technique to the branching program, recursively, with parameter b . For simplicity we first assume that S is a power of b . Each application of the technique reduces the path length by a factor b . Hence with $\log_b S$ applications we can reduce the path length to 1.

In one application, we have an \exists quantifier over $b - 1$ nodes, corresponding to an Or gate with fan-in S^{b-1} , and then a \forall quantifier over b smaller paths, corresponding to an And gate with fan-in b . This gives a tree with $S^{b-1} \cdot b \leq S^b$ leaves. Iterating, the number of leaves will be

$$(S^b)^{\log_b S}.$$

Each leaf can be connected to the input bit on which it depends. The size of the circuit is at most c times the number of leaves.

If S is not a power of b we can view the branching program as having $S' \leq bS$ nodes where S' is a power of b . **QED**

The following is a uniform version of Theorem 7.10, and the proof is similar. It shows that we can speed up space-bounded computation with alternations.

Theorem 7.11. Any $f \in L$ is also in $\Sigma_{c_f/\epsilon}\text{Time}(n^\epsilon)$, for any $\epsilon > 0$.

Proof. Let G be the configuration graph of f . Note this graph has n^{c_f} nodes. We need to decide if the start configuration reaches the accept configuration in this graph within $t := n^{c_f}$ steps.

We apply to this graph the checkpoint technique recursively, with parameter $b := n^{\epsilon/2}$. Each application of the technique reduces the path length by a factor b . Hence with c_f/ϵ applications we can reduce the path length to

$$\frac{t}{b^{c_f/\epsilon}} = \frac{n^{c_f}}{n^{c_f/2}} \leq 1.$$

Each quantifier ranges over $b \log n^{c_f} = c_f n^{\epsilon/2} \log n \leq n^\epsilon$ bits for large enough n .

There remains to check a path of length 1, i.e., an edge. The endpoints of this edge are two configurations u and v which depend on the quantified bits. The machine can compute the two endpoints in time $\log^c m$ where m is the total number of quantified bits, using rapid access. Once it has u and v it can check if u leads to v in one step by reading one bit from the input. Note $m \leq n^\epsilon \cdot c_f/\epsilon$, so $\log^c m \leq c_f n^\epsilon$. **QED**

We note that by the generic simulation of alternating time by small-depth circuits in Lemma 6.2, the above theorem also gives a result similar to Theorem 7.10.

7.4 Reductions: L vs. P

Again, we can use reductions to related the space complexity of problems. In particular we can identify the problems in P which have space-efficient algorithms iff every problem in P does.

Definition 7.6. A problem f is P-complete if $f \in P$ and $f \in L \Rightarrow L = P$.

Definition 7.7. The circuit value problem: Given a circuit C and an input x , compute $C(x)$.

Theorem 7.12. Circuit value is P-complete.

Proof. Circuit value is in P since we can evaluate one gate at the time. Now let $g \in P$. We can reduce computing g on input x to a circuit value instance, as in the simulation of TMs by circuits in Theorem 2.5. The important point is that this reduction is computable in L. Indeed, given an index to a gate in the circuit, we can compute the type of the gate and index to its children via simple arithmetic, which is in L by Theorem 7.5, and some computation which only depends on the description of the P-time machine for g . **QED**

Definition 7.8. The monotone circuit value problem: Given a circuit C with no negations and an input x , compute $C(x)$.

Exercise 7.9. Prove that monotone circuit value is P-complete.

Recall from section 7.2.3 that finding solutions to linear systems

$$Ax = b$$

has space-efficient algorithms. Interestingly, if we generalize equalities to inequalities the problem becomes P complete. This set of results illustrates a sense in which “linear algebra” is easier than “optimization.”

Definition 7.9. The linear inequalities problem: Given a $d \times d$ matrix A of integers and a d -dimensional vector, determine if the system $Ax \leq b$ has a solution over the reals.

Theorem 7.13. Linear inequalities is P-complete.

Proof. The ellipsoid algorithm shows that Linear inequalities is in P, but we will not discuss this classic result.

Instead, we focus on showing how given a circuit C and an input x we can construct a set of inequalities that are satisfiable iff $C(x) = 1$.

We shall have as many variables v_i as gates in the circuit, counting input gates as well.

For an input gate $g_i = x_i$ add equation $v_i = x_i$.

For a Not gate $g_i = \text{Not}(g_j)$ add equation $v_i = 1 - v_j$.

For an And gate $g_i = \text{And}(g_j, g_k)$ add equations $0 \leq v_i \leq 1, v_i \leq v_j, v_i \leq v_k, v_j + v_k - 1 \leq v_i$.

The case of Or is similar, or can be dispensed by writing an Or using Not and And.

Finally, if g_i is the output gate add equation $v_i = 1$.

We claim that in every solution to $Av \leq b$ the value of v_i is the value in [2] of gate g_i on input x . This can be proved by induction. For input and Not gates this is immediate. For an And gate, note that if $v_j = 0$ then $v_i = 0$ as well because of the equations $v_i \geq 0$ and $v_i \leq v_j$. The same holds if $v_k = 0$. If both v_j and v_k are 1 then v_i is 1 as well because of the equations $v_i \leq 1$ and $v_j + v_k - 1 \leq v_i$. **QED**

7.4.1 Nondeterministic space

Because of the insight we gained from considering non-deterministic time-bounded computation in section §5.1, we are naturally interested in non-deterministic space-bounded computation. In fact, perhaps we will gain even more insight, because this notion will really challenge our understanding of computation.

For starters, let us define non-deterministic space-bounded computation. A naive approach is to define it using the quantifiers from section §6.4, leading to the class $\exists \cdot L$. This is an ill-fated choice:

Exercise 7.10. Prove $\exists \cdot L = \exists \cdot P$.

Instead, non-deterministic space is defined in terms of non-deterministic TMs.

Definition 7.10. A function $f : [2]^* \rightarrow [2]$ is computable in $\text{NSpace}(s(n))$ if there is a two-tape TM which on input x never writes on the first tape and never uses more than $s(n)$ cells on the second, and moreover:

1. The machine is equipped with a special “Guess” state. Upon entering this state, a *guess bit* is written on the work tape under the head.
2. $f(x) = 1$ iff there exists a choice for the guess bits that causes the machine to output 1.

We define

$$\text{NL} := \bigcup_d \text{NSpace}(d \log n),$$

$$\text{NPSPACE} := \bigcup_d \text{NSpace}(n^d).$$

How can we exploit this non-determinism? Recall from section 7.2.2 that reachability in *undirected* graphs is in L. It is unknown if the same holds for directed graphs. However, we can solve it in NL.

Definition 7.11. The directed reachability problem: Given a directed graph G and two nodes s and t , decide if there is a path from s to t .

Theorem 7.14. Directed reachability is in NL.

Proof. The proof simply amounts to guessing a path in the graph. The algorithm is as follows:

“On input G, s, t , let $v := s$.

For $i = 0$ to $|G|$:

 If $v = t$, accept.

 Guess a neighbor w of v . Let $v := w$.

If you haven’t accepted, reject.”

The space needed is $|v| + |i| = c \log |G|$. **QED**

We can define NL completeness similarly to NP and P completeness, and have the following result.

Theorem 7.15. Directed reachability is NL-complete. That is, it is in NL and it is in L iff $L = \text{NL}$.

Exercise 7.11. Prove this.

Recall by definition $\text{Space}(s(n)) \subseteq \text{NSpace}(s(n))$. We showed $\text{Space}(s(n)) \subseteq \text{Time}(n^c c^{s(n)})$ in Theorem 7.1. We can strengthen the inclusion to show that it holds even for non-deterministic space.

Theorem 7.16. $\text{NSpace}(s(n)) \subseteq \text{Time}(n^c c^{s(n)})$.

Proof. On input x , we compute the configuration graph G of M on input x . The nodes are the configurations, and there is an edge from u to v if the machine can go from u to v in one step. Then we solve reachability on this graph in power time, using say breadth-first-search.

QED

The next theorem shows that non-deterministic space is not much more powerful than deterministic space: it buys at most a square. Contrast this with the P vs. NP question! The best deterministic simulation of NP that we know is the trivial $\text{NP} \subseteq \text{Exp}$. Thus the situation for space is entirely different.

How can this be possible? The high-level idea, which was used already in Lemma 7.1, can be cast as follows:

Unlike time, space can be reused.

Theorem 7.17. $\text{NSpace}(s) \subseteq \text{Space}(cs^2)$, for every function $s = s(n) \geq \log n$. In particular, $\text{NPSpace} = \text{PSpace}$.

Proof. We use the checkpoint technique with parameter $b = 2$, and re-use the space to verify the smaller paths. Let N be a non-deterministic TM computing a function in $\text{NSpace}(s(n))$. We aim to construct a deterministic TM M that on input x returns

$$\text{Reach}(C_{\text{start}}, C_{\text{accept}}, c^{s(n)}),$$

where $\text{Reach}(u, v, t)$ decides if v is reachable from u in $\leq t$ steps in the configuration graph of N on input x , and C_{start} is the start configuration, C_{accept} is the accept configuration, and $c^{s(n)}$ is the number of configurations of N .

The key point is how to implement Reach .

Computing $\text{Reach}(u, v, t)$
 For all “middle” configurations m
 If both $\text{Reach}(u, m, t/2) = 1$ and $\text{Reach}(m, v, t/2) = 1$ then Accept.
 Reject

Let $S(t)$ denote the space needed for computing $\text{Reach}(u, v, t)$. We have

$$S(t) \leq cs(n) + S(t/2).$$

This is because we can re-use the space for two calls to Reach . Therefore, the space for $\text{Reach}(C_{\text{start}}, C_{\text{accept}}, c^{s(n)})$ is

$$\leq cs(n) + cs(n) + \dots + cs(n) \leq cs^2(n).$$

QED

To set the stage for the next result, recall that we do not know if $\text{Ntime}(t)$ is closed under complement. It is generally believed not to be, and we showed that if it is then the PH collapses Exercise 6.3.

What about space? Theorem 7.17 shows $\text{NSpace}(s) \subseteq \text{Space}(cs^2)$. Because the latter is closed under complement, up to a quadratic loss in space, non-deterministic space is closed under complement.

Can we avoid squaring the space?

Yes! This is weird!

Theorem 7.18. The complement of Path is in NL. In particular, NL is closed under complement.

Proof. We want a non-deterministic 2TM that given G, s , and t accepts if there is no path from s to t in G .

For starters, suppose the machine has computed the number C of nodes reachable from s . *The key idea is that there is no path from s to t iff there are C nodes different from t reachable from s .* Thus, knowing C we can solve the problem as follows

Algorithm for deciding if there is no path from s to t , given C :

Initialize Count=0; Enumerate over all nodes $v \neq t$
 Guess a path from s of length $|G|$. If path reaches v , increase Count by 1
 If Count = C Accept, else Reject.

There remains to compute C .

Let A_i be the nodes at distance $\leq i$ from s , and let $C_i := |A_i|$. Note $A_0 = \{s\}$, $c_0 = 1$. We seek to compute $C = C_n$.

To compute C_{i+1} from C_i , enumerate nodes v (candidate in A_{i+1}). For each v , enumerate over all nodes w in A_i , and check if $w \rightarrow v$ is an edge. If so, increase C_{i+1} by 1.

The enumeration over A_i is done guessing C_i nodes and paths from s . If we don't find C_i nodes, we reject. **QED**

Exercise 7.12. Given a graph G and nodes s, t show how to compute in L a graph G' and nodes s', t' s.t. there is a path from s to t in G iff there is no path from s' to t' in G' , and $|G'| \leq |G|^c$.

7.5 TiSp

So far in this chapter we have focused on bounding the space usage. For this, the TM model was sufficient, as remarked at the beginning. It is natural to consider algorithms that operate in little time *and* space. For this, of course, whether we use TMs or RAMs makes a difference.

Definition 7.12. Let $\text{TiSp}(t, s)$ be the functions computable on a RAM that on every input $x \in [2]^n$ runs in time $t(n)$ and only uses memory cells 0 to $s(n) - 1$.

Exercise 7.13. Prove $L = \bigcup_d \text{TiSp}(n^d, d)$.

An alternative definition of TiSp would allow the RAM to access $s(n)$ cells anywhere in memory. One can maintain a data structure to show that this alternative definition is equivalent to Definition 7.12.

To illustrate the relationship between TiSp , Time, and Space, consider undirected reachability. It is solvable in $\text{Time}(n \log^c n)$ by breadth-first search, and in logarithmic space by Theorem 7.8. But it isn't known if it is in $\text{TiSp}(n \log^a n, a \log n)$ for some constant a .

Exercise 7.14. Prove the following version of Theorem 7.11: $\text{TiSp}(n^a, n^{1-\alpha}) \subseteq \Sigma_{ca/\alpha} \text{Time}(n)$ for any $a \geq 1$ and $\alpha > 0$.

The following is a non-uniform version of TiSp .

Definition 7.13. A branching program of length t and width W is a branching program where the nodes are partitioned in t layers L_1, L_2, \dots, L_t where nodes in L_i only lead to nodes in L_{i+1} , and $|L_i| \leq W$ for every i .

Thus t represents the time of the computation, and $\log W$ the space.

Recall that Theorem 7.10 gives bounds of the form $\geq cn^2 / \log n$ on the size of branching program (without distinguishing between length and width). For branching programs of length t and width W this bound gives $t \geq cn^2 / W \log n$. Note this gives nothing for power width like $W = n^2$. The state-of-the-art for power width is $t \geq \Omega(n \sqrt{\log n} / \log \log n)$ (in fact the bound holds even for subexponential width).

With these definitions in hand we can refine the connection between branching programs and small-depth circuits in Theorem 7.10 for circuits of depth 3.

Theorem 7.19. Let $f : [2]^n \rightarrow [2]$ be computable by a branching program with width W and time t . Then f is computable by an alternating depth-3 circuit with $\leq 2^{c\sqrt{t \log W}}$ wires.

We will later show explicit functions that require depth-3 circuits of size $2^{c\sqrt{n}}$. Theorem 7.19 shows that improving this would also improve results for small-width branching programs, a refinement of the message emphasized after Theorem 7.10.

A more general version of Theorem 7.19. states that for any parameter b one can have a depth-3 circuit with

$$2^{b \log W + t/b + \log t}$$

wires, output fan-in W^b , and input fan-in t/b . Interestingly, this tradeoff essentially matches known impossibility results for depth-3 circuits!

Exercise 7.15. Prove Theorem 7.19.

7.6 Problems

Problem 7.1. Consider the class $\Sigma_{a(n)}\text{Time}(t(n))$ where the number of alternations is $a(n)$ on inputs of length n (as opposed to being fixed to i for every input as in $\Sigma_i\text{Time}(t(n))$).

Prove $L \neq \Sigma_{a(n)}\text{Time}(n^\delta)$ for any growing $a(n)$ and any constant δ . (Hint: Use a slight extension of the previous problem.)

Problem 7.2. In this problem we will explore a beautiful illustration of the power of L. The Simplify problem is defined as follows. The input is a string over the alphabet $\{u, v, u^{-1}, v^{-1}\}$. Such a string can be simplified by removing adjacent pairs of the type $uu^{-1}, u^{-1}u, vv^{-1}, v^{-1}v$. The problem asks whether a given input string becomes the empty string after simplification.

For example, $uv^{-1}vu^{-1}$ simplifies to uu^{-1} and then to the empty string. On the other hand, the string $uv^{-1}u^{-1}$ cannot be simplified to the empty string.

We will show that Simplify is in L.

In this problem $|x|$ denotes absolute value.

(1) For integers i consider the matrices

$$U_i := \begin{bmatrix} 1 & 2i \\ 0 & 1 \end{bmatrix}, V_i := \begin{bmatrix} 1 & 0 \\ 2i & 1 \end{bmatrix}.$$

Show that $U_i U_j = U_{i+j}$ and so in particular $U_i^{-1} = U_{-i}$; and show the same for the V_i .

(2) Let $\begin{bmatrix} x \\ y \end{bmatrix}$ be a vector and let $i \neq 0$. Show that if $|x| < |y|$ then $\begin{bmatrix} x' \\ y' \end{bmatrix} := U_i \begin{bmatrix} x \\ y \end{bmatrix}$ has $|x'| > |y'|$. Conversely, show that if $|x| > |y|$ then $\begin{bmatrix} x' \\ y' \end{bmatrix} := V_i \begin{bmatrix} x \\ y \end{bmatrix}$ has $|x'| < |y'|$. (This is the so-called ping-pong lemma.)

(3) Show that an alternating product of matrices $U_{i_1} V_{i_2} U_{i_3} V_{i_4} \cdots U_{i_k}$ where the i_j are not zero is not equal to the identity matrix. Note that we begin and end with a U matrix, and we alternate between U and V .

(4) Show that a product of matrices $U_{i_1} V_{i_2} U_{i_3} \cdots V_{i_k}$ where the i_j are not zero is not equal to the identity matrix. Note that we begin with U but end with a V matrix, and as before we alternate U and V matrices. (Hint: Reduce to (3) by multiplying on the left by M^{-1} and on the right by M .)

(5) Show that Simplify is in L.

7.7 Notes

Theorem 7.2 is from [94].

Theorem 7.4: [140].

Theorem 7.8 is from [162]. This came quite as a shock during my own PhD, because the proof is very simple. Later a simpler yet proof has appeared [165], using even more standard machinery.

Theorem 7.9 follows from [52] which in fact establishes a stronger result, specifically it shows that the problem is in NC^2 , a class we encounter in Chapter 9.

The use of CRR for arithmetic is from [31], which also contains several reductions among arithmetical problems. Some of the steps are from the earlier work [133]. For a discussion of the complexity of division, see [9].

For a compendium of P-complete problems see [75].

Theorem 7.11 goes back to [141].

As for NP, a compendium of problems complete for P is available [75].

Theorem 7.17: [167].

Theorem 7.18 was obtained independently in [97, 189]. The proof in [97] uses a logical formalism, the proof we presented is closer to the one in [189]. This question is a.k.a. the second LBA problem, from [117]. As usual, had the answer been different, it would have had applications to the first LBA problem, which is the basically question whether Theorem 7.17 is tight.

See [6, 33] for the state-of-the-art bounds for power-width branching programs.

Chapter 8

Three impossibility results for 3Sat

We should turn back to a traditional separation technique – diagonalization.

In this chapter we put together many of the techniques we have seen to obtain several impossibility results for 3Sat. The template of all these results (and others, like those mentioned in section §5.1) is similar. All these results prove time bounds of the form $t \geq n^{1+\alpha}$ where $\alpha \in (0, 1)$. One can optimize the methods to push α close to 1, but even establishing $\alpha = 1$ seems out of reach, and there are known barriers for current techniques.

8.1 Impossibility I

We begin with the following remarkable result.

Theorem 8.1. Either $3\text{Sat} \notin \text{L}$ or $3\text{Sat} \notin \text{Time}(n^{1+\epsilon})$ for some constant ϵ .

Note that we don't know if $3\text{Sat} \in \text{L}$ or if $3\text{Sat} \in \text{Time}(n \log^{10} n)$. In particular, Theorem 8.1 implies that any algorithm for 3Sat either must use super-logarithmic space or time $n^{1+\epsilon}$.

Proof. We assume that what we want to prove is not true and derive the following striking contradiction with the hierarchy Theorem 3.4:

$$\begin{aligned} \text{Time}(n^2) &\subseteq \text{L} \\ &\subseteq \bigcup_d \Sigma_d \text{Time}(n) \\ &\subseteq \text{Time}(n^{1.9}). \end{aligned}$$

The first inclusion holds by the assumption that $3\text{Sat} \in \text{L}$ and the fact that any function in $\text{Time}(n^2)$ can be reduced to 3Sat in log-space, by Theorem 5.1 and the discussion after that.

The second inclusion is Theorem 7.11.

For the third inclusion, the assumption that $3\text{Sat} \in \text{Time}(n^{1+\epsilon})$ for every ϵ implies that $\text{NTime}(dn) \subseteq \text{Time}(n^{1+\epsilon})$ for every d and ϵ , by the quasi-linear-time completeness of 3Sat , Theorem 5.4. Now apply Exercise 6.2. **QED**

8.2 Impossibility II

We now state and prove a closely related result for TiSp . We seek to rule out algorithms for 3Sat that simultaneously use little space and time, whereas in Theorem 8.1 we even ruled out the possibility that there are two distinct algorithms, one optimizing space and the other time. The main gain is that we will be able to handle much larger space: power rather than log.

Theorem 8.2. $3\text{Sat} \notin \text{TiSp}(n^{1+c_\epsilon}, n^{1-\epsilon})$, for any $\epsilon > 0$.

The important aspect of Theorem 8.1 is that it applies to the RAM model; stronger results can be shown for space-bounded TMs.

Exercise 8.1. Prove that $\text{Palindromes} \notin \text{TM-TiSp}(n^{1+c_\epsilon}, n^{1-\epsilon})$, for any $\epsilon > 0$. ($\text{TM-TiSp}(t, s)$ is defined as $\text{Space}(s)$, cf. Definition 7.1, but moreover the machine runs in at most t steps.) Hint: This problem has a simple solution. Give a suitable simulation of TM-TiSp by 1TM , then apply Theorem 3.1.

Proof. We assume that what we want to prove is not true and derive the following contradiction with the hierarchy Theorem 3.4:

$$\begin{aligned} \text{Time}(n^{1+\epsilon}) &\subseteq \text{TiSp}(cn^{(1+\epsilon)(1+c_\epsilon)}, cn^{(1+\epsilon)(1-\epsilon)}) \\ &\subseteq \text{TiSp}(n^{1+c_\epsilon}, cn^{1-\epsilon^2}) \\ &\subseteq \Sigma_{c_\epsilon} \text{Time}(n) \\ &\subseteq \text{Time}(n^{1+\epsilon/2}). \end{aligned}$$

The first inclusion holds by the assumption, padding, and the fact that 3Sat is complete under reductions s.t. each bit is computable in time (and hence space) $n^{o(1)}$, a fact we do not prove here. **QED**

Exercise 8.2. Finish the proof by justifying the remaining inclusions.

8.3 Impossibility III

So far our impossibility results required bounds on space. We now state and prove a result that applies to time. Of course, as discussed in Chapter 3, we don't know how to prove that, say, 3Sat cannot be computed in linear time on a 2TM . For single-tape machines, we can

prove quadratic bounds, for palindromes (Theorem 3.1) and 3Sat (Problem 4.2). Next we consider an interesting model which is between 1TM and 2TM and is a good indication of the state of our knowledge.

Definition 8.1. A 1.5TM is like a 2TM except that the input tape is read-only.

Theorem 8.3. 3Sat requires time n^{1+c} on a 1.5TM.

Exercise 8.3. Prove Theorem 8.3 following this guideline:

1. Let M be a 1.5TM running in time $t(n)$. Divide the read-write tape of M into consecutive blocks of b cells, shifted by an offset $i < b$. (So the the first cells of the blocks include $i, i + b, i + 2b, \dots$) Prove that for every input $x \in [2]^n$ there is i such that the sum of the lengths of the crossing sequences between any adjacent blocks of the computation M on x is at most $t(n)/b$. Here a crossing sequence also encodes the position of the head on the input tape, and the time at which each crossing occurs.
2. Prove that $1.5\text{TM-Time}(n^{1.1}) \subseteq \exists y \in [2]^{n^{1-c}} \text{TiSp}(n^c, n^{1-c})$. (The right-hand side is the class of functions $f : [2]^* \rightarrow [2]$ for which there is a RAM M that on input (x, y) , where $|x| = n$, runs in time n^c and uses memory cells $0..n^{1-c}$ and s.t. $f(x) = 1 \Leftrightarrow \exists y \in [2]^{n^{1-c}} M(x, y) = 1$.)
3. Conclude the proof.

Notes

The introductory quote is from [58], where Theorem 8.1 is proved. This influential work ignited a whole research area. For a survey (not up to date) see [200] or [220]. For the limitations of this type of results, see [43]. Theorem 8.3 is proved in [130] for NTime instead of 3Sat. The result for 3Sat appeared in [201].

Chapter 9

Log-depth circuits

In this chapter we investigate circuits of logarithmic depth. Again, several surprises lay ahead, including a solution to the teaser in Chapter 1!

Let us begin slowly with some basic properties of these circuits so as to get familiar with them. The next exercise shows that circuits of depth $d \log n$ for a constant d also have power size, so we don't need to bound the size separately.

Exercise 9.1. A circuit of depth d has size $\leq c^d$ without loss of generality.

The next exercises shows how to compute several simple functions by log-depth circuits.

Exercise 9.2. Prove that the Or, And, and Parity functions on n bits have circuits of depth $c \log n$.

Prove that any $f : [2]^n \rightarrow [2]$ computable by an AC of depth d and size $s \geq n$ is also computable by a circuit of depth $cd \log s$ and size s^c .

Next, let us relate these circuits to branching programs. The upshot is that circuits of logarithmic depth are a special case of power-size branching programs, and the latter are a special case of circuits of log-square depth.

Theorem 9.1. Directed reachability has circuits of depth $c \log^2 n$ and size n^c . In particular, the same holds for any function in NL, and any function with power-size branching programs.

Proof. On input a graph G on u nodes and two nodes s and t , let M be the $u \times u$ transition matrix corresponding to G , where $M_{i,j} = 1$ iff edge $j \rightarrow i$ is in G .

Transition matrices are multiplied as normal matrices, except that “+” is replaced with “ \vee ,” which suffices to know connectivity. To answer directed reachability we compute entry t of $M^u v$, where v has a 1 corresponding to s and 0 everywhere else. (We can modify the graph to add a self-loop on node t so that we can reach t in exactly u steps iff we reach t in any number of steps.)

Computing M^u can be done by squaring $c \log u$ times M . Each squaring can be done in depth $c \log u$, by Exercise 9.2. This establishes the first claim, since $u \leq n$.

The “in particular” follows because those functions can be reduced to directed reachability efficiently. **QED**

Conversely, we have the following.

Theorem 9.2. Any function $f : [2]^n \rightarrow [2]$ computed by a circuit of depth d can be computed by a branching program of size 2^d .

In particular, functions computed by circuits of logarithmic depth can be computed by branching programs of power size.

Later in this chapter we will prove a stronger and much less obvious result.

Proof. We proceed by induction on the depth of the circuit C . If the depth is 1 then C is either a constant or an input bit, and a branching program of size 1 is available by definition.

Suppose the circuit C has the form $C_1 \wedge C_2$. By induction, C_1 and C_2 have branching programs B_1 and B_2 each of size 2^{d-1} . A branching program B for C of size 2^d is obtained by rewiring the edges leading to states labelled 1 in B_1 to the start state of B_2 . The start state of B is the start state of B_1 . **QED**

Exercise 9.3. Finish the proof by analyzing the case $C = C_1 \vee C_2$.

Definition 9.1. NC^i is the class of functions $f : [2]^* \rightarrow [2]^*$ computable by circuits that have depth $a \log^i n$ and size n^a , for some constant a . The circuits are uniform if they can be computed in L.

The class NC^0 is also of great interest. It can be more simply defined as the class of functions where each output bit depends on a constant number of input bits. We will see many surprising useful things that can be computed in this class.

The previous results give, for uniform circuits:

$$\text{NC}^0 \subseteq \text{NC}^1 \subseteq \text{L} \subseteq \text{NL} \subseteq \text{NC}^2 \subseteq \text{NC}^3 \subseteq \dots \subseteq \text{P}.$$

The only inclusion known to be strict is the first one:

Exercise 9.4. Prove that $\text{NC}^0 \neq \text{NC}^1$. (Mostly to practice definitions.)

9.1 The power of NC^1 : Arithmetic

In this section we illustrate the power of NC^1 by showing that the same basic arithmetic which we saw is doable in L (Theorem 7.5) can in fact be done in NC^1 as well.

Theorem 9.3. The following problems are in NC^1 :

1. Addition of two input integers.
2. Iterated addition: Addition of any number of input integers.

3. Multiplication of two input integers.
4. Iterated multiplication: Multiplication of any number of input integers.
5. Division of two integers.

Exercise 9.5. Prove Item 1. in Theorem 9.3.

Iterated addition is surprisingly non-trivial. We can't use the methods from the proof of Theorem 7.5. Instead, we rely on a new and very clever technique.

Proof of Item 2. in Theorem 9.3.. We use "2-out-of-3:" Given 3 integers X, Y, Z , we compute 2 integers A, B such that

$$X + Y + Z = A + B,$$

where each bit of A and B only depends on three bits, one from X , one from Y , and one from Z . Thus A and B can be computed in NC^0 .

If we can do this, then to compute iterated addition we construct a tree of logarithmic depth to reduce the original sum to a sum 2 terms, which we add as in Item 1.

Here's how it works. Note $X_i + Y_i + Z_i \leq 3$. We let A_i be the least significant bit of this sum, and B_{i+1} the most significant one. Note that A_i is the XOR $X_i + Y_i + Z_i$, while B_{i+1} is the majority of X_i, Y_i, Z_i . **QED**

The following corollary will also be used to solve the teaser in Chapter 1.

Corollary 9.1. Majority is in NC^1 .

Exercise 9.6. Prove it.

Exercise 9.7. Prove Item 3. in Theorem 9.3.

Next we turn to iterated multiplication. The idea is to follow the proof for L in section 7.2.1. We shall use CRR again. The problem is that we still had to perform iterated multiplication, albeit only in \mathbb{Z}_p for $p \leq n^c$. One more mathematical result is useful now:

Theorem 9.4. If p is a prime then $(\mathbb{Z}_p - \{0\})$ is a cyclic group, meaning that there exists a generator $g \in (\mathbb{Z}_p - \{0\}) : \forall x \in (\mathbb{Z}_p - \{0\}), x = g^i$, for some $i \in \mathbb{Z}$.

Example 9.1. For $p = 5$ we can take $g = 2$: $2^0 = 1, 2^1 = 2, 2^2 = 4, 2^3 = 8 = 3$.

Proof of Item 4. in Theorem 9.3. We follow the proof for L in section 7.2.1. To compute iterated product of integers r_1, r_2, \dots, r_t modulo p , use Theorem 9.4 to compute exponents e_1, e_2, \dots, e_t s.t.

$$r_i = g^{e_i}.$$

Then $\prod_i r_i \text{ mod } p = g^{\sum_i e_i}$. We can use Item 2. to compute the iterated addition of the exponents. Note that computing the exponent of a number mod p , and vice versa, can be done in log-depth since the numbers have $c \log n$ bits (as follows for example by combining Theorem 2.4 and Exercise 9.2). **QED**

One can also compute division, and make all these circuits uniform, but we won't prove this now.

9.2 Computing with 3 bits of memory

We now move to a surprising result that in particular strengthens Theorem 9.2. For a moment, let's forget about circuits, branching programs, etc. and instead consider a new, minimalistic type of programs. We will have 3 one-bit registers: R_0, R_1, R_2 , operating modulo 2. We allow the following operations

$$\begin{aligned} R_i+ &= R_j, \\ R_i+ &= R_j x_k \end{aligned}$$

where x_k is an input bit, for any $i, j \in \{0, 1, 2\}$, with $i \neq j$. (Talk about RISC!) Here $R_i+ = R_j$ means to add the content of R_j to R_i , while $R_i+ = R_j x_k$ means to add $R_j x_k$ to R_i , where $R_j x_k$ is the product (a.k.a. And) of R_j and x_k .

Definition 9.2. For i, j and $f : [2]^n \rightarrow [2]$ we say that a program is *for* (or *equivalent to*)

$$R_i+ = R_j f$$

if for every input x and initial values of the registers, executing the program is equivalent to the instruction $R_i+ = R_j f(x)$, where note that R_j and R_k are unchanged.

Also note that if we repeat twice a program for $R_i+ = R_j f$ then no register changes (recall the sum is modulo 2, so $1 + 1 = 0$). This feature is critically exploited later to “clean up” computation.

We now state and prove the surprising result. It is convenient to state it for circuits with Xor instead of Or gates. This is without loss of generality since $x \vee y = x + y + x \wedge y$.

Theorem 9.5. Suppose $f : [2]^n \rightarrow [2]$ is computable by circuits of depth d with Xor and And gates. For every $i \neq j$ there is a program of length $\leq c4^d$ for

$$R_i+ = R_j f.$$

Once such a program is available, we can start with register values $(0, 1, 0)$ and $i = 0, j = 1$ to obtain $f(x)$ in R_0 .

Proof. We proceed by induction on d . When $d = 1$ the circuit is simply outputting a constant or one of the input bits, which we can compute with the corresponding instructions. (If the circuit is the constant zero then the empty program would do.)

Proceeding with the induction step:

A program for $R_i+ = R_j(f_1 + f_2)$ is simply given by the concatenation of (the programs for)

$$\begin{aligned} R_i+ &= R_j f_1 \\ R_i+ &= R_j f_2. \end{aligned}$$

Less obviously, a program for $R_i+ = R_j(f_1 \wedge f_2)$ is given by

$$\begin{aligned}
R_i+ &= R_k f_1 \\
R_k+ &= R_j f_2 \\
R_i+ &= R_k f_1 \\
R_k+ &= R_j f_2.
\end{aligned}$$

QED

Exercise 9.8. Prove that the program for $f_1 \wedge f_2$ in the proof works. Write down the contents of the registers after each instruction in the program.

A similar proof works over other fields as well.

We can now address the teaser Theorem 1.1 from Chapter 1.

Proof of Theorem 1.1. Combine Corollary 9.1 with Theorem 9.5. **QED**

Corollary 9.2. Iterated product of 3x3 matrices over \mathbb{F}_2 is complete for NC^1 under projections.

That is, the problem is in NC^1 and for any $f \in \text{NC}^1$ and n one can write a sequence of $t = n^c$ 3x3 matrices M_1, M_2, \dots, M_t where each entry is either a constant or an input variable x_i s.t. for every $x \in [2]^n$:

$$\prod_{i=1}^t M_i \cdot \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} = \begin{bmatrix} f(x) \\ 1 \\ 0 \end{bmatrix}.$$

Exercise 9.9. Prove this.

Recall from Chapter 7 (see in particular section 7.4.1) that various graph reachability problems are complete for space-bounded computation. In particular, one can reduce any function computable by branching programs of size s to iterated multiplication of $s \times s$ matrices over \mathbb{F}_2 .

Exercise 9.10. Prove this. Specifically, prove that any function $f : [2]^n \rightarrow [2]$ computable by branching programs of size s can be reduced via projections to iterated multiplication of $s \times s$ matrices over \mathbb{F}_2 . Does a similar result hold for NL?

Hence major open questions about computation are related to the following “purely mathematical question” that doesn’t make any direct reference to computation:

Can (any one entry of) the product of $s \times s$ matrices be reduced via projection to the product of $s^d 3 \times 3$ matrices, for some constant d ? (That is, the former product has s^3 variables x_i and each entry in the latter product is either a variable x_i or 0 or 1, all of this is over \mathbb{F}_2 .)

Specifically, if the answer is positive then $\text{L} \subseteq \text{NC}^1$. And if not then we have an explicit problem that is not in NC^1 . To obtain an equivalence one can consider the class $\oplus\text{L}$, see the notes.

9.3 Linear-size log-depth

It is unknown whether NP has linear-size circuits of logarithmic depth! But there is a non-trivial simulation of such circuits by AltCkts of depth 3 of sub-exponential size.

Theorem 9.6. Any circuit $C : [2]^n \rightarrow [2]$ of size an and depth $a \log n$ has an equivalent AC of depth 3 and size $2^{c_{an}/\log \log n}$.

The idea is... yes! Once again, we are going to guess computation. It is possible to guess the values of about $n/\log \log n$ wires to reduce the depth to say $0.1 \log n$, and the rest is brute-forced.

Notes

One can verify that the matrices in 9.2 form a group, so we are representing computation via group products. The landmark paper on such representations is [135], although the general idea is actually already in [116]. The critical insight in the former is that this representation is efficient for small-depth circuits. After [135], several related simulations were discovered, such as [35]. Corollary 9.2 is from [135]. Theorem 9.5 is a variant of this result from [35]. The proof we presented follows [50].

Iterated multiplication of matrices over \mathbb{F}_2 is complete [54] for the class $\oplus L$. Thus the “purely mathematical question” is equivalent to (the non-uniform version of) the question $\oplus L = NC^1$.

Theorem 9.6 is from [196]. For an exposition, see [207].

Chapter 10

Constant-depth circuits

In this section we further investigate circuits of constant depth, focusing on two pervasive classes, which indeed we have already encountered under different names.

10.1 Threshold circuits

Definition 10.1. A *threshold circuit*, abbreviated TC, is a circuit made of Majority gates (of unbounded fan-in). We also denote by TC the class of functions f computable by a TC of depth d and size n^d for some constant d .

Exercise 10.1. Prove that $AC \subseteq TC \subseteq NC^1$.

TCs are one of the frontiers of our knowledge. It isn't known how to prove impossibility results even for TCs of depth 3 and size, say, n^2 . As usual, a good explanation for this ignorance is the power of TCs, of which we give several examples next.

Exercise 10.2. A function $f : [2]^* \rightarrow [2]$ is *symmetric* if it only depends on the weight of the input. Prove that any symmetric function is in TC. Hint: Similar to a part of Exercise 6.13. You may want to start by proving that deciding if the input weight is equal to a fixed value is in TC.

The result $PH \subseteq \text{Maj} \cdot \text{Maj} \cdot P$ obtained in 6.13 in particular yields the following.

Theorem 10.1. Any function f in AC has TCs of depth 3 and size $2^{\log^c f n}$.

Exercise 10.3. Prove Theorem 10.1 but for depth 4 instead of 3, using only Theorem 6.5 from Chapter 6. Hint: Use Exercise 10.2.

Theorem 10.2. The following problems are in TC:

1. Addition of two input integers.
2. Iterated addition: Addition of any number of input integers.

3. Multiplication of two input integers.
4. Iterated multiplication: Multiplication of any number of input integers.
5. Division of two integers.

The proof follows closely that for NC^1 in section §9.1 (which in turn was based on that for L). Only iterated addition requires a new idea.

Exercise 10.4. Prove the claim about iterated addition. (Hint: Write input as $n \times n$ matrix, one number per row. Divide columns into blocks of $t = c \log n$.)

10.2 TC vs. NC^1

Another great question is whether $\text{TC} = \text{NC}^1$. For any d , we can show that functions in NC^1 , such as Parity, require depth- d TCs of size $\geq n^{1+c \log d}$, and this is tight up to constants. A natural question is whether we can prove stronger bounds for harder functions in NC^1 . A natural candidate is iterated multiplication of 3×3 matrices. The following result shows that, in fact, stronger bounds would already prove “the whole thing,” that is, $\text{TC} \neq \text{NC}^1$.

Theorem 10.3. Let G be the set of 3×3 matrices of \mathbb{F}_2 . Suppose that the product of n elements in G can be computed by TCs of size n^k and depth d . Then for any ϵ the product can also be computed by TCs of size $d' n^{1+\epsilon}$ and depth $d' := cdk \log 1/\epsilon$.

The same result applies to any constant-size group G – we state it for matrices for concreteness.

Proof. Exploiting the associativity of the problem, we compute the product recursively according to a regular tree. The root is defined to have level 0. At Level i we compute n_i products of $(n^{1+\epsilon}/n_i)^{1/k}$ matrices. At the root ($i = 0$) we have $n_0 = 1$.

By the assumption, each product at Level i has TCs of size $n^{1+\epsilon}/n_i$ and depth d . Hence Level i can be computed by TCs of size $n^{1+\epsilon}$ and depth d .

We have the recursion

$$n_{i+1} = n_i \cdot (n^{1+\epsilon}/n_i)^{1/k}.$$

The solution to this recursion is $n_i = n^{(1+\epsilon)(1-(1-1/k)^i)}$, see below.

For $i = ck \log(1/\epsilon)$ we have $n_i = n^{(1+\epsilon)(1-\epsilon^2)} > n$; this means that we can compute a product of $\geq n$ matrices, as required.

Hence the total depth of the circuit is $d \cdot ck \log(1/\epsilon)$, and the total size is the depth times $n^{1+\epsilon}$.

It remains to solve the recurrence. Letting $a_i := \log_n n_i$ we have the following recurrence for the exponents of n_i .

$$\begin{aligned} a_0 &= 0 \\ a_{i+1} &= a_i(1 - 1/k) + (1 + \epsilon)/k = a_i\beta + \gamma \end{aligned}$$

where $\beta := (1 - 1/k)$ and $\gamma := (1 + \epsilon)/k$.

This gives

$$a_i = \gamma \sum_{j \leq i} \beta^j = \gamma \frac{1 - \beta^{i+1}}{1 - \beta} = (1 + \epsilon)(1 - \beta^{i+1}).$$

QED

Were the recursion of the form $a'_{i+1} = a'_i + (1 + \epsilon)/k$ then obviously a'_{ck} would already be $\geq 1 + \epsilon$. Instead for a_i we need to get to $ck \log(1/\epsilon)$.

10.3 Impossibility results for AC

In this section we prove impossibility results for ACs, matching several settings of parameters mentioned earlier (cf. section §7.3).

To set the stage, let's prove strong results for depth 2, that is, DNFs.

Exercise 10.5. Prove that Majority requires DNFs of size $\geq 2^{cn}$. Hint: What if you have a term with $< n/2$ variables?

As discussed, 2^{cn} bounds even for depth 3 ACs are unknown, and would imply major separations. The following is close to the state-of-the-art for depth d .

Theorem 10.4. Let C be an AC of depth d and size s computing Majority on n bits. Then $\log^d s \geq c\sqrt{n}$.

Recall from section §7.3 that a stronger bound for an explicit function would have major consequences; in particular the function cannot be in L.

10.3.1 Impossibility results by polynomial method (a.k.a. low-degree approximation)

The proof uses the simulation of circuits by low-degree polynomials which we saw in Theorem 6.5. Specifically, we use the following corollary:

Corollary 10.1. Let $C : [2]^n \rightarrow [2]$ be an alternating circuit of depth d and size s . Then there is a polynomial p over \mathbb{F}_2 of degree $\log^d s/\epsilon$ such that $\mathbb{P}_x[C(x) \neq p(x)] \leq \epsilon$.

Proof. Theorem 6.5 gave a distribution P on polynomials s.t. for every x we have

$$\mathbb{P}_P[C(x) \neq P(x)] \leq \epsilon.$$

Averaging over x we also have

$$\mathbb{P}_{x,P}[C(x) \neq P(x)] \leq \epsilon.$$

Hence we can fix a particular polynomial p s.t. the probability over x is $\leq \epsilon$, yielding the result. **QED**

We then show that Majority cannot be approximated by such low-degree polynomials. The key result is the following:

Lemma 10.1. Every function $f : [2]^n \rightarrow [2]$ can be written as $f(x) = p_0(x) + p_1(x) \cdot \text{Maj}(x)$, for some polynomials p_0 and p_1 of degree $\leq n/2$. This holds for every odd n .

Proof. Let M_0 be the set of strings with weight $\leq n/2$. We claim that for every function $f : M_0 \rightarrow [2]$ there is a polynomial p_0 of degree $\leq n/2$ s.t. p_0 and f agree on M_0 .

To verify this, consider the monomials of degree $\leq n/2$. We claim that (the vectors corresponding to) their truth tables over M_0 are linearly independent. This means that any polynomial gives a different function over M_0 , and because the number of polynomials is the same as the number of functions, the result follows. **QED**

Exercise 10.6. Prove the claim in the proof.

Proof of Theorem 10.4. Apply Corollary 10.1 with $\epsilon = 1/10$ to obtain p . Let S be the set of inputs on which $p(x) = C(x)$. By Lemma 10.1, any function $f : S \rightarrow [2]$ can be written as

$$f(x) = p_0(x) + p_1(x) \cdot p(x).$$

The right-hand side is a polynomial of degree $\leq d' := n/2 + \log^d(cs)$. The number of such polynomials is the number of possible choices for each monomial of degree i , for any i up to the degree. This number is

$$\prod_{i=0}^{d'} 2^{\binom{n}{i}} = 2^{\sum_{i=0}^{d'} \binom{n}{i}}.$$

On the other hand, the number of possible functions $f : S \rightarrow [2]$ is

$$2^{|S|}.$$

Since a polynomial computes at most one function, taking logs we have

$$|S| \leq \sum_i^{d'} \binom{n}{i}.$$

The right-hand side is at most $2^n(1/2 + c \log^d(s)/\sqrt{n})$, since each binomial coefficient is $\leq c2^n/\sqrt{n}$.

On the other hand, $|S| \geq 0.9 \cdot 2^n$.

Combining this we get

$$0.9 \cdot 2^n \leq 2^n(1/2 + c \log^d(s)/\sqrt{n}).$$

This implies

$$0.4 \leq c \log^d(s) / \sqrt{n},$$

proving the theorem. **QED**

Exercise 10.7. Explain why Theorem 10.4 holds even if the circuits have Parity gates (in addition to Or and And gates).

If as in the above exercise we allow for parity gates, Theorem 10.4 is the strongest known bound. Stronger bounds are only known for functions computable in classes related to exponential time.

However, for AC a sharper technique is known that allows us to replace the \sqrt{n} in Theorem 10.4 with n for several functions such as parity. This is presented next.

10.3.2 Switching lemmas

A *restriction* ρ is an assignment of the variables to $\{0, 1, \star\}$, i.e., some variables are replaced with constants, while those assigned to *star* \star are left “alive.” In a *random restriction*, the stars are selected uniformly at random, and also each unrestricted variable is set to a uniform bit. We denote by N_s the number of restrictions with exactly s stars.

Exercise 10.8. $N_s = 2^{n-s} \binom{n}{s}$.

For a restriction ρ with s stars and $f : [2]^n \rightarrow [2]$ we denote by $f_\rho : [2]^s \rightarrow [2]$ the restricted function. The switching lemma shows that important classes of functions simplify dramatically when “hit” by a random restriction.

Definition 10.2. A *decision tree* of depth d is a branching program where every path has length $\leq d$.

In particular, depth- d decision trees compute 2^d -local functions, cf. Definition 4.3.

Lemma 10.2. Let $C : [2]^n \rightarrow [2]$ be the Or of functions f_i each on w bits. Let ρ be a random restriction with s stars. The probability that f_ρ is not a decision tree of depth d is $\leq (cws/n)^d$.

To illustrate, f can be a DNF or a CNF with terms of size $w = c \log n$. The lemma says that if we pick a uniform restriction with $s := \sqrt{n}$ stars, and set $d = 10$, then the restricted function on s bits is a decision tree of depth d except with probability $((c \log n)/n)^{10} \leq 1/n^c$. Note in particular the restricted function depends only on $2^d = 1024$ bits, even though the original function may depend on all the bits. If we take larger d , the error probability gets even smaller, so small in fact that we can take a union bound and simultaneously simplify many DNFs. Doing this several times allows us to collapse an AC to a low-depth decision tree. We state and prove this consequence next, trading simplicity of exposition for parameter optimization.

Corollary 10.2. Let $C : [2]^n \rightarrow [2]$ be an AC of depth d size s . Let ρ be a random restriction with n^{c_d} stars. The probability that C_ρ is not a decision tree of depth $\log s$ is $\leq s2^{-n^{c_d}}$.

Proof. Set $w := \log s$. We view ρ as successive applications of restrictions whose number of stars is square root of the number of variables. View the circuit as having depth $d+1$ and the input gates have fan-in $1 \leq w$. The first application of Lemma 10.2 gives error $\leq (cw/\sqrt{n})^w$. In the good case, the input gates now are decision trees of depth $\leq w$. We can write this as a CNF or DNF with terms of size $\leq w$, and merge the output gate with the gates in the next level in the circuit, which are now computing Ors (or Ands) of functions on $\leq w$ bits. The next application of Lemma 10.2 gives error $\leq (cw/n^{1/4})^w$, and so on. **QED**

One can use the switching lemma to prove exponential lower bounds to compute explicit functions by small-depth ACs. The simplest example is parity, given next. In this case, we also prove an exponentially strong correlation bound.

Definition 10.3. The *correlation* between two functions $f, g : [2]^n \rightarrow [2]$ is defined as

$$|\mathbb{E}_x[(-1)^{f(x)+g(x)}]| = |\mathbb{P}[f(x) = g(x)] - \mathbb{P}[f(x) \neq g(x)]| = |1 - 2\mathbb{P}[f(x) \neq g(x)]|.$$

We also introduce the notation $\mathbb{E}e[X]$ for $\mathbb{E}[(-1)^X]$ which allows us to write correlation as $|\mathbb{E}e[f(x) + g(x)]|$.

Thus the correlation between two functions is a measure of how often the functions agree on a uniform input. To illustrate parameters, if $f = g$ or $f = 1 - g$ then the correlation is one. If f and g disagree on exactly one input the correlation is $1 - 1/2^n - 1/2^n = 1 - 2/2^n$. Also, for any f , most functions g have correlation close to 0. (As can be proved via counting arguments as in Chapter 3.)

Corollary 10.3. The correlation between parity and an AC of depth d and size s is $\leq s2^{-n^{c_d}}$.

Proof. The correlation between parity on m bits and decision trees of depth $< m$ is zero. View a uniform input as first picking a restriction, and then filling the stars. By Corollary 10.2, after picking the restriction the circuit is a decision tree of depth $\log s$, which is strictly less than the number of remaining stars, except with probability $s2^{-n^{c_d}}$. **QED**

Exercise 10.9. State and prove via a reduction from Corollary 10.3 an impossibility result for Majority. Does a strong correlation bound hold as well?

10.3.3 Proof of Lemma 10.2

The simplest case: Or of n bits Here f is simply the Or of n bits x_1, x_2, \dots, x_n . In the restriction some of the bits may become 0, others 1, and others yet may remain unfixed, i.e., assigned to stars. Those that become 0 you can ignore, while if some become 1 then the whole circuit C becomes 1.

We will show that the number of restrictions for which the restricted circuit $C|_\rho$ requires decision trees of depth $\geq d$ is small.

For this simple case, a straightforward proof of a stronger bound exists.

Exercise 10.10. Give it. For concreteness, think $d = 10$ and $s = \sqrt{n}$.

We give an alternative argument which we can then extend to the general case. We are going to encode (or map) such restrictions using (or: to a restriction)... with no stars (that is, just a 0/1 assignment to the variables). The gain is clear: just think of a restriction with zero stars versus a restriction with one star. Recall Exercise 10.8. We have $N_0 = 2^n$, while $N_1 = 2^{n-1} \cdot n$, so $N_0/N_1 \leq c/n$. Note that this an upper bound on the error probability.

A critical observation is that we only want to encode restrictions for which $C|_\rho$ requires large depth. So ρ does not map any variable to 1, for else the Or is 1 which has decision trees of depth 0.

The way we are going to encode ρ is this: *Simply replace the stars with ones.* To go back, replace the ones with stars. We are using the ones in the encoding to “signal” where the stars are.

Hence, the number of bad restrictions is at most $N_0 = 2^n$, which is tiny compared to the number $N_s = \binom{n}{s} 2^{n-s}$ of restrictions with s stars (cf. Exercise 10.8). The error probability is then

$$\frac{2^n}{\binom{n}{s} 2^{n-s}} = \frac{2^s}{\binom{n}{s}} \leq (2s/n)^s.$$

Here we used $\binom{n}{s} \geq (n/s)^s$. This is stronger than Lemma 10.2. (We can assume $d \leq s$, since every function on s bits has decision trees of depth s , and so for $d \geq s$ the error probability is 0.)

The medium case: Or of functions on disjoint inputs So, again, let’s take a random restriction ρ with exactly s stars. Some of the functions may become 0, others 1, and others yet may remain unfixed. Those that become 0 you can ignore, while if some become 1 then the whole circuit becomes 1.

As before, we will show that the number of restrictions for which the restricted circuit $C|_\rho$ requires decision trees of depth $\geq d$ is small. To accomplish this, we are going to encode/map such restrictions using/to a restriction with just $s - d$ stars, plus a little more information. As we saw already, the gain in reducing the number of stars is clear. In particular, standard calculations show that saving d stars reduces the number of restrictions by a factor $(cs/n)^d$.

Exercise 10.11. Prove $N_{s-d}/N_s \leq (cs/n)^d$. Guideline: Recall Exercise 10.8. To make the calculations as simple as possible, set $s := \sqrt{n}$ and consider the following way of selecting s stars among n variables: First pick $s - d$ stars. Then divide the remaining $n - (s - d)$ variables into d blocks of $(n - (s - d))/d \geq cn/d$ variables. Pick exactly one star in each block. Conclude the proof.

The auxiliary information will give us a factor of w^d , leading to the claimed bound. Specifically, as before, recall that we only want to encode restrictions for which $C|_\rho$ requires large depth. So no function in $C|_\rho$ is 1, for else the circuit is 1 and has decision trees of depth 0. Also, you have d stars among inputs to functions that are unfixed (i.e., not even fixed to 0), for else again you can compute the function reading less than d bits. Because

the functions are unfixed, there is a setting for those d stars (and possibly a few more stars – that would only help the argument) that make the corresponding functions 1. We are going to pick precisely that setting in our restriction ρ' with $s - d$ stars. This allows us to “signal” which functions had inputs with the stars we are saving (namely, those that are the constant 1). To completely recover ρ , we add extra information to indicate where the stars were. The saving here is that we only have to say where the stars are among w symbols, not n .

Specifically, we can encode the positions of the stars with an element $a \in [cw]^d$, indicating which of the w symbols is a star, and also whether the star is the last in this function. To recover the restriction, we look for the first function that’s set to 1. We then read a to find out how many stars were there and their positions. Then we move to the second function that’s fixed to 1. Again we look at a to know how many stars were there and what their positions were, and so on.

The general case: Or of functions on any subset of w bits ...isn’t really different. First, the number of functions does not play a role, so you can think you have functions on any possible subset of w bits, where some functions may be constant. The idea is the same, except we have to be slightly more careful because when we set values for the stars in one function we may also affect other functions. The idea is to fix one function at the time. Specifically, starting with ρ , consider the first function f that’s not made constant by ρ . So the inputs to f have some stars. As before, let us replace the stars with constants that make the function f equal to the constant 1, and append the extra information that allows us to recover where these stars were in ρ .

We’d like to repeat the argument. Note however we only have guarantees about $C|_\rho$, not $C|_\rho$ with some stars replaced with constants that make f equal to 1. We also can’t just jump to the 2nd function that’s not constant in $C|_\rho$, since the “signal” fixing for that might clash with the fixing for the first – this is where the overlap in inputs makes things slightly more involved. Instead, because $C|_\rho$ required decision tree depth at least d , we note there have to be some assignments to the m stars in the input to f so that the resulting, further restricted circuit still requires decision tree depth $\geq d - m$ (else $C|_\rho$ has decision trees of depth $< d$). We append this assignment to the auxiliary information and we continue the argument using the further restricted circuit.

10.3.4 The original switching lemma

We now state and prove an earlier, weaker switching lemma, whose proof is however simpler and more intuitive.

Lemma 10.3. Let $f : [2]^n \rightarrow [2]$ be a k -CNF. Let ρ be a random restriction with $\mathbb{P}[*] = 1/n^c$. The probability that f_ρ is not c_k -local is $\leq 1/n^k$.

Proof. Induction on k . For $k = 1$, f is an And. If the And is on $\geq ck \log n$ bits then f_ρ will be constant with the desired probability. If the And is on $\leq ck \log n$ bits then the prob. that

f_ρ depends on $\geq c_k$ bits is

$$\leq \binom{ck \log n}{c_k} n^{-c \cdot c_k} \leq 1/n^{c_k}. \quad (10.1)$$

For the induction step, suppose there are $\geq c_k \log n$ Or gates with disjoint inputs. Since each Or gate is 0 after the restriction w.p. $\geq 1/c_k$, the result follows.

Otherwise, there is a set C of $\leq c_k \log n$ variables that touches every Or gate. For every assignment to these variables, we can apply the induction hypothesis, and do a union bound. Also, like in equation (10.1), the prob. that $\geq c_k$ variables in C are set to $*$ by ρ is $\leq 1/n^{c_k}$.
QED

10.4 ACs can sample

We showed in the earlier section that ACs cannot compute parity and majority. However, ACs can sample input-output pairs of these functions. In other words, this is a problem that ACs can't solve, but for which nevertheless they can create instances together with their solutions. These results have applications and again point to the unsuspected power of these circuits. Throughout we assume that the inputs to the circuits is uniform.

Exercise 10.12. Give a 2-local map $C : [2]^n \rightarrow [2]^{n+1}$ whose output distribution is $(X, \text{parity}(X))$ for uniform $X \in [2]^n$.

Sampling $(X, \text{majority}(X))$ by ACs is more involved and beautiful, and is only known to be possible up to a small error.

Exercise 10.13. Try to sample $(X, \text{majority}(X))$ by ACs for a few minutes. Write down what you tried.

The first step is sampling a uniform permutation.

Lemma 10.4. There is an AC whose output distribution is 2^{-n^c} close to a uniform permutation π over $[n]$, represented as n blocks of $c \log n$ bits where block i has $\pi(i)$ in binary.

Exercise 10.14. Assume Lemma 10.4. For any $i \leq n$ give an AC whose output distribution is 2^{-n^c} -close to a uniform n -bit strings of weight i . Give an AC whose output distribution is 2^{-n^c} -close to $(X, \text{majority}(X))$ for uniform $X \in [2]^n$.

Proof of Lemma 10.4. The main technique is known as “dart throwing:” we view the input random bits as random pointers p_1, p_2, \dots, p_n into $m \gg n$ cells. We then write i in the p_i -th cell (empty cells get “*”). If there are no collisions, the ordering of $[n]$ in the cells gives a random permutation of $[n]$. However, it is not clear how to explicitly write out this permutation using small depth, because to determine the image of i one needs to count how many cells before p_i are occupied, which cannot be done in small depth.

The key insight is to view the cells as representing the permutation in a different format, one from which we can explicitly write out the permutation in small depth. The format is

known as the canonical form for the cyclic notation. We now briefly review it. Just like the standard format, the alternative format represents a permutation via an array $A[1..n]$ whose entries contain all the elements $[n]$. However, rather than thinking of $A[i]$ as the image of i , we think of the entries of A as listing the cycles of the permutation in order. Each cycle is listed starting with its smallest element, and cycles are listed in decreasing order of the first element in the cycle. This format allows for computing the permutation efficiently: the image of i is the element to the right of i in A , unless the latter element is the beginning of a new cycle, in which case the image of i is the first element in the cycle containing i . Identifying the first element of a cycle is easy, because it is smaller than any element preceding it in A . The benefit of this format is that it works even if the array A has $m \gg n$ cells, of which $m - n$ are empty and marked by “*.”

One can now verify that computing the image of i can be done in AC. Here in particular we use the fact that such circuits can, given an array A and an index i , compute the least $j > i$ such that $A[j]$ is not “*”. This can be accomplished by trying all j in parallel, noting that one can determine if a fixed j is the least $j > i$ such that $A[j]$ is not “*” using one unbounded fan-in And.

To conclude the proof of the lemma, generate ℓ uniform and independent sets of pointers p_1^i, \dots, p_n^i , $i = 1, \dots, \ell$, where each pointer has range $[m]$ for m the smallest power of 2 larger than $2n^2$ (thus each pointer can be specified by $\log m$ bits).

If there exists i such that the pointers p_1^i, \dots, p_n^i are all distinct (i.e., there are no collisions), then run the above algorithm on the output corresponding to the first such i . This results in a random permutation.

Since the pointers are chosen independently, the probability that there is no such i is

$$\Pr[\forall i \exists j, k \leq n : p_j^i = p_k^i] = \Pr[\exists j, k \leq n : p_j^1 = p_k^1]^\ell \leq (1/2)^\ell.$$

Choosing $\ell := n$ proves the lemma. **QED**

10.5 ACC

We denote AC augmented with gates computing $\text{mod } m$ by $\text{AC}[m]$. ACC (alternating circuits with counters) refers to any m . We also denote by $\text{AC}[m]$ the class of functions computable by $\text{AC}[m]$ circuits of size n^d and depth d for some constant d , and similarly for ACC.

Techniques based on polynomials, such as those discussed in section 10.3.1 are effective to prove impossibility results against $\text{AC}[m]$ if m is prime. Exercise 10.7 refers to the fundamental case $m = 2$. But they break down when m is composite. It is consistent with our knowledge that any function in Exp has ACC of depth c and size n^c with $\text{mod } 6$ gates. It is open if Majority has such circuits.

For any m , $\text{AC}[m]$ can be simulated by polynomials. The simulation is incomparable to the one we saw previously for special cases of m such as $m = 2$. In the new simulation We work with polynomials over the integers and then map the output to a boolean value.

Equivalently, we can think of the polynomial as a depth-2 circuit. On the other hand, this new simulation works for every input as opposed to most inputs.

Lemma 10.5. Any $AC[d]$ of size n^d and depth d has an equivalent depth-2 circuit of size $2^{\log^{c_d} n}$ where the output is a symmetric function (i.e., only depends on the number of bits that are 1 in the input to that gate) and the other gates are And with fan-in $\leq \log^{c_d} n$.

As far as we know, general circuits are equivalent to ACC! Yet there is one thing that we can say about functions computable in ACC that we don't know for PCKt. We can solve ACC-Sat better than brute-force search. Using this, and Lemma 10.5, and diagonalization, one can prove the following result, which we don't know how to prove in other ways.

Theorem 10.5. $NExp \neq ACC$.

The technique involve guessing circuits so it does not seem applicable to functions in NP.

10.6 Notes

Theorem 10.1: [8].

The impossibility result for parity in TC is from [100].

Theorem 10.3: [10, 47].

Impossibility results for AC are among the most famous results in complexity and were obtained in the 80's in [61, 5, 223, 158, 179, 85] via various techniques. The first two works obtained super-power results, the others exponential. Each work gives a negative result for a symmetric function and hence applies to majority as well. The proof of Theorem 10.4 we gave is from [180], though that paper splashes mathematics in a way that doesn't help me (Hilbert function, Grobner basis, "in the projective case the ideal $I(S)$ is *homogeneous* (or *graded*) if..." etc.).

Stronger bounds for AC with parity gates for less explicit functions are in [203].

Lemma 10.4 is from [132, 80]. Our presentation follows [211].

Lemma 10.5 is from [225, 34].

Theorem 10.5 is from [221]. One step of the original proof was somewhat indirect and it was streamlined later in [104].

10.6.1 Myth creation: The switching lemma

I must admit I had a good run (private communication)

The history of science is littered with anecdotes about misplaced credit. Because it does not matter if it was A or B who did it; it only matters if it was I or not I. The only point of this section is the disproportionate amount of credit that the work has received within and without our community (typically due to inertia and snowball effects rather than malice). Of course, at some level this doesn't matter. You can call Chebichev's polynomials rainbow

sprinkles and the math doesn't change. And yet at some other level maybe it does matter a little, for science isn't yet a purely robotic activity.

Random restrictions have been used in complexity theory since at least the 60's [186]. The first dramatic use in the context of AC0 is due to [61, 5]. These works proved a *switching lemma* the amazing fact that a DNF gets simplified by a random restriction to the point that it can be written as a CNF, so you can collapse layers and induct. (An exposition is given below.) Using it, they proved super-polynomial lower bounds for AC0. The proof in [61], presented in section 10.3.4, is very nice and if I want to get a quick intuition of why switching is at all possible, I often go back to it. [5] is also a brilliant paper, and long, unavailable online for free, filled with a logical notation which makes some people twitch. The first symbol of the title says it all, and may be the most obscene ever chosen:

$$\Sigma_1^1.$$

Subsequently, [223] proved exponential lower bounds of the form 2^{n^c} , with a refined analysis of the switching lemma. The bounds are tight, except for the constant c which depends on the depth of the circuit. Finally, the star of this section [84, 85] obtained $c = 1/(\text{depth} - 1)$.

Yao's paper doesn't quite state that a DNF can be written exactly as a CNF, but it states that it can be approximated. Hastad's work is the first to prove that a DNF can be written as a CNF, and in this sense his statement is cleaner than Yao's. However, Yao's paper states explicitly that a small circuit, after being hit by a restriction, can be set to constant by fixing few more bits.

The modern formulation of the switching lemma says that a DNF can be written as a *shallow decision tree* (and hence a small CNF). This formulation in terms of decision trees is actually not explicit in Hastad's work. Beame, in his primer [30], credits Cai with this idea and mentions several researchers noted Hastad's proof works in this way.

Another switching lemma trivia is that the proof in Hastad's thesis is actually due to Boppana; Hastad's original argument -- of which apparently no written record exists -- was closer to Razborov's later proof.

So, let's recap. Random restrictions are already in [186]. The idea of switching is already in [61, 5]. You already had three analyses of these ideas, two giving superpolynomial lower bounds and one [223] giving exponential. The formulation in terms of decision trees isn't in [85], and the proof that appears in [85] is due to Boppana.

Still, I would guess [85] is more well known than all the other works above combined. [223] did have a following at the time -- I think it appeared in the pop news. But hey -- have you ever heard of Yao's switching lemma?

The current citation counts offer mixed support for my thesis:

FSS: 1351

Y: 732

H - paper "Almost optimal..." 867

H - thesis: 582

But it is very hard to use citation information. The two H citations overlap, and papers are cited for various reasons. For example FSS got a ton of citations for the connection to oracles (which has nothing to do with switching lemmas).

Instead it's instructive to note the type of citations that you can find in the literature:

Hastad's switching lemma is a cornerstone of circuit complexity [No mention of FSS, A, Y]

Hastad's Switching Lemma is one of the gems of computational complexity [Notes below in passing it builds on FSS, A, Y]

The wikipedia entry is also telling:

In computational complexity theory, Hastad's switching lemma is a key tool for proving lower bounds on the size of constant-depth Boolean circuits. Using the switching lemma, Johan Haastad (1987) showed that... [No mention of FSS,A,Y]

I think that 99% of the contribution of this line of research is the *amazing idea* that random restrictions simplify a DNF so that you can write it as a CNF and collapse. 90% of the rest is analyzing this to get superpolynomial lower bounds. And 90% of whatever is left is analyzing this to get exponential lower bounds.


Going back to something I mentioned in the first post, I want to emphasize that Hastad during talks makes a point of reminding the audience that the idea of random restrictions is due to Sipser, and of Boppana's contribution. And I also would like to thank him for his help with this post.

OK -- so maybe this is so, but it must then be the case that [85] is the final word on this stuff, like the ultimate tightest analysis that kills the problem. Actually, it is not tight in some regimes of interest, and several cool works of past and recent times address that. In the end, I can only think of one reason why [85] entered the mythology in ways that other works did not, the reason that I carefully sidestepped while composing this post: *â*.

Perhaps one reason behind the aura of the switching lemma is that it's hard to find examples. It would be nice to read: If you have this extreme DNF here's what happens, on the other hand for this other extreme DNF here's what happens, and in general this always works and here's the switching lemma. *Examples are forever* – Erdos. Instead the switching lemma is typically presented as *blam!*: an example-free encoding argument which feels *deus ex machina*, as in this crisp presentation by Thapen. I have tried to give a slightly different exposition of the encoding argument.

Chapter 11

Proofs



The notion of proof is pervasive. We have seen many proofs in this book until now. But the notion extends to others realms of knowledge, including empirical science, law, and more. Complexity theory has contributed a great deal to the notion of proof, with important applications in several areas such as cryptography.

11.1 Static proofs

As remarked in section 5.1.1, we can think of problems in NP as those admitting a solution that can be verified efficiently, namely in P. Let us repeat the definition of NP using the suggestive letter V for verifier.

Definition 11.1. A function $f : X \subseteq [2]^* \rightarrow [2]$ is in NP iff there is $V \in P$ (called “verifier”) and $d \in \mathbb{N}$ s.t.:

$$f(x) = 1 \Leftrightarrow \exists y \in [2]^{|x|^d} : V(x, y) = 1.$$

We are naturally interested in fast proof verification, and especially the complexity of V . It turns out that proofs can be encoded in a format that allows for very efficient verification. This message is already in the following.

Theorem 11.1. For any input length n , V in Definition 11.1 can be taken to be a 3CNF of size n^d .

That is, whereas when defining NP as a proof system we considered arbitrary verifiers V in P, in fact the definition is unchanged if one selects a very restricted class of verifiers: small 3CNFs.

Proof. This is just a restatement of Theorem 5.1. **QED**

This extreme reduction in the verifier’s complexity is possible because we are allowing proofs to be long, longer than the original verifier’s running time. If we don’t allow for that, such a reduction is not known. Such “bounded proofs” are very interesting to study, but we shall not do so now.

Instead, we ask for more. The 3CNF in the above theorem still depends on the entire proof. We can ask for a verifier that only depends on few bits of the proof. Taking this to the extreme, we can ask whether V can only read a constant number of bits from y . Without randomness, this is impossible.

Exercise 11.1. Suppose V in Definition 11.1 only reads $\leq d$ bits of y , for a constant d . Show that the corresponding class would be the same as P.

Surprisingly, if we allow randomness this is possible. Moreover, the use of randomness is fairly limited – only logarithmically many bits – yielding the following central characterization.

Theorem 11.2. A function $f : X \subseteq [2]^* \rightarrow [2]$ is in NP iff there is $V \in P$ and $d \in \mathbb{N}$ s.t.:

$$\begin{aligned} f(x) = 1 &\Rightarrow \exists y \in [2]^{|x|^d} : \mathbb{P}_{r \in [2]^{d \log |x|}} [V(x, y, r) = 1] = 1, \\ f(x) = 0 &\Rightarrow \forall y \in [2]^{|x|^d} : \mathbb{P}_{r \in [2]^{d \log |x|}} [V(x, y, r) = 1] < 0.01, \\ &\text{and moreover } V \text{ reads } \leq d \text{ bits of } y. \end{aligned}$$

Exercise 11.2. Prove the “only if” in Theorem 11.2 in the specific case $f = 0.01\text{-Gap-3Sat}$.

Given this exercise, the “only if” direction for any problem in NP follows from the advanced result that any problem in NP can be map reduced to 0.01-Gap-3Sat (which is essentially Theorem 4.10, except we did not claim map reductions or a specific constant there).

Exercise 11.3. Prove the “if” in Theorem 11.2.

11.2 Proofs that yield nothing but their validity: Zero-knowledge

In Theorem 11.2 the verifier gains “constant confidence” about the validity of the proof, just by inspecting a constant number of bits. Hence the verifier “learns” at most a constant number of bits of the proof. This is remarkable, but we can further ask if we can modify

the proof so that the verifier “learns nothing” about the proof. Such proofs are called *zero knowledge* and are extensively studied and applied.

We sketch how this is done for Gap-3Color, which is also NP-complete. Rather than a single proof y , now the verifier will receive a random proof Y . This Y is obtained from a 3 coloring y by randomly permuting colors (so for any y the corresponding Y is uniform over 6 colorings). The verifier will pick a random edge and inspect the corresponding endpoints, and accept if they are different.

The verifier learn nothing because all that they see is two random different color. One can formalize “learning nothing” by noting that the verifier can produce this distribution by themselves, without looking at the proof. (So why does the verifier gain anything from y ? The fact that a proof y has been written down means that colors have been picked so that every two endpoints are uniform colors, something that the verifier is not easily able to reproduce.)

This gives a zero-knowledge proof for verifiers that follow the protocol of just inspecting an edge. In a cryptographic setting one has to worry about verifiers which don’t follow the protocol. Using cryptographic assumptions, one can force the verifiers to follow the protocol by considering an *interactive* proof: First a proof y is committed to but not revealed, then the verifier selects an edge to inspect, and only then the corresponding colors are revealed, and only those. This protocol lends itself to a physical implementation that can astonish the right audiences.

11.3 Interactive proofs

We now consider interactive proofs. Here the verifier V engages in a protocol with a prover P . Given an input x to both V and P , the verifier asks questions, the prover replies, the verifier asks more questions, and so on. The case of NP corresponds to the prover simply sending y to V .

It turns out that it suffices for the verifier to send uniformly random strings Q_1, Q_2, \dots bits to P . This leads to a simple definition.

Definition 11.2. A function $f : X \subseteq [2]^* \rightarrow [2]$ admits an efficient interactive proof, abbreviated IP, if there is $V \in \mathcal{P}$ and $d \in \mathbb{N}$ such that for every $x \in [2]^n$, letting $b := n^d$:

- If $f(x) = 1$ then $\exists P : [2]^* \rightarrow [2]^b$ such that

$$V(P(Q_1), P(Q_1, Q_2), \dots, P(Q_1, Q_2, \dots, Q_b)) = 1$$

for every $Q_1, Q_2, \dots, Q_b \in [2]^b$.

- If $f(x) = 0$ then $\forall P : [2]^* \rightarrow [2]^b$ we have

$$\mathbb{P}_{Q_1, Q_2, \dots, Q_b \in [2]^b} [V(P(Q_1), P(Q_1, Q_2), \dots, P(Q_1, Q_2, \dots, Q_b)) = 1] \leq 1/3.$$

The following amazing result shows the power of interactive proofs, compared to non-interactive proofs. We can think of NP as “reading a book” and IP as “going to class and asking questions.” We don’t yet know how to replace teachers with books.

Theorem 11.3. $\text{IP} = \text{PSPACE}$.

In the rest of this section we present the main ideas in the proof of 11.3, establishing a weaker result. In particular we show that IP contains problems not known to be in NP.

Theorem 11.4. Given a field \mathbb{F} , an arithmetic circuit $C(x_1, x_2, \dots, x_v)$ over \mathbb{F} computing a polynomial of degree d , and an element $s \in \mathbb{F}$, deciding if

$$\sum_{x_1, x_2, \dots, x_v \in [2]} C(x_1, x_2, \dots, x_v) = s \quad (11.1)$$

is in IP, whenever $(1 - d/q)^v \geq 2/3$.

Proof. If $v = 1$ then V can decide this question by itself, by evaluating the circuit. For larger v we give a way to reduce v by 1.

As the first prover answer, V expects a polynomial p of degree d in the variable x , which is meant to be

$$s'(x) := \sum_{x_2, x_3, \dots, x_n \in [2]} C(x, x_2, x_3, \dots, x_n).$$

V checks if $p(0) + p(1) = s$, and if not rejects. Otherwise, it recursively runs the protocol to verify that

$$\sum_{x_2, x_3, \dots, x_n \in [2]} C(Q_1, x_2, x_3, \dots, x_n) = p(Q_1). \quad (11.2)$$

This concludes the description of the protocol. We now verify its correctness.

In case equation (11.1) is true, P can send polynomials that cause V to accept.

In case equation (11.1) is false, $s'(0) + s'(1) \neq s$. Hence, unless V rejects right away because $p(0) + p(1) \neq s$, we have $p \neq s'$. The polynomials p and s' have degree $\leq d$. Hence by Lemma 2.3

$$\mathbb{P}_{Q_1}[p(Q_1) \neq s'(Q_1)] \geq 1 - d/q.$$

When this event occurs, equation (11.2) is again false, and we can repeat the argument. Overall, the probability that we maintain a false statement throughout the protocol is $\geq (1 - d/q)^v$. **QED**

Corollary 11.1. Given a 3CNF formula ϕ and $k \in \mathbb{N}$, deciding if ϕ has exactly k satisfying assignments is in IP.

The proof uses a far-reaching technique: *arithmetization*. We construct an arithmetic circuit C_ϕ over a field \mathbb{F} which agrees with ϕ on *boolean* inputs, but that can then be evaluated over other elements of the field.

Exercise 11.4. Prove Corollary 11.1.

The study of interactive proofs is rich. Many aspects are of interest, including:

- The efficiency of the prover (does it have to be unbounded, randomized, etc.), and of the verifier.
- The number of rounds.
- The tradeoff between the error and the other parameters.

11.4 Interactive proofs for muggles

tbd

11.5 Notes

Interactive proofs were put forth simultaneously in [21, 73]. The latter paper also introduced zero-knowledge. The zero-knowledge protocol for 3Color is from [70].

Theorem 11.3 is from [128, 170], with the last paper proving it as stated. For the history of this most famous result see [18].

“Muggles” are from [72].

For more on interactive proofs and zero knowledge see the recent book [190].

Chapter 12

Data structures

Data structures aim to maintain data in memory so as to be able to support various operations, such as answering queries about the data, and updating the data. The study of data structures is fundamental and extensive. We distinguish and study in turn two types of data structure problems: *static* and *dynamic*. In the former the input is given once and cannot be modified by the queries. In the latter queries can modify the input; this includes classical problems such as supporting insert, search, and delete of keys.

12.1 Static data structures

Definition 12.1. A static data-structure problem is simply a function $f : [2]^n \rightarrow [q]^m$. A data structure for f with space s , word size w and time t is a way to write f as

$$f(x) = h(g(x))$$

where $g : [2]^n \rightarrow [2^w]^s$, $h : [2^w]^s \rightarrow [q]^m$, and each output bit of h depends on $\leq t$ input words (we think of s as divided into words of length w).

Here we have n bits of input data about which we would like to answer m queries. Often the queries and or the word size are boolean, i.e., $2^w = q = 2$. Another typical setting is $q = 2^w$. The data structure aims to accomplish this by storing the input into s bits of memory. This map is arbitrary, with no bound on resources. But after that, each query can be answered very fast, by reading only t words. In general, these words can be read adaptively. But for simplicity we focus on the case in which the locations are fixed by the data structure and the same for every input $x \in [2]^n$.

Exercise 12.1. Consider the data structure problem $f : [2]^n \rightarrow [2]^m$ where $m = n^2$ and query $(i, j) \in \{1, 2, \dots, n\}^2$ is the parity of the input bits from i to j .

Give a data structure for this problem with $s = n$, $w = 1$, and $t = 2$.

Exercise 12.2. Show that any data-structure problem $f : [2]^n \rightarrow [2]^m$ has a data structure with $w = 1$ and the following parameters:

- (1) $s = m$ and $t = 1$, and
- (2) $s = n$ and $t = n$.

Exercise 12.3. Prove that for every n and $m \leq 2^{n/2}$ there exist functions $f : [2]^n \rightarrow [2]^m$ s.t. any data structure with space $s = m/2$ and $w = 1$ requires time $t \geq n - c$.

By contrast, next we present the the best known impossibility result.

Definition 12.2. A function $f : [2]^n \rightarrow [q]^m$ is d -wise uniform, or simply d -uniform, if any d output coordinates are uniform when the input to f is uniform.

Theorem 12.1. Let $f : [q]^d \rightarrow [q]^q$ be d -wise uniform. Let q be a power of 2 and $c \log q \leq d \leq q^c$. Then any data structure with $w = \log q$ using space s (which recall is measured in words of w bits) and time t has:

$$t \geq c \frac{\log q}{\log(s/d)}.$$

Interpreting the input as coefficients of a degree $d - 1$ univariate polynomial over \mathbb{F}_q and outputting its evaluations shows that such functions exists, and are in P. Below we give a surprising data structure that nearly matches the theorem.

To match previous parameters note that $n = d \log q = dw$, and $m = \log q$. Hence the bound is $t \geq c(\log m)/\log(sw/n)$. Note that sw is the space of the data structure measured in bits. It follows that if sw is linear in n then $t \geq c \log m$. This result remains non-trivial for s slightly super-linear. But remarkably, if $sw = n^{1+c}$ then nothing is known (for m power in n one only gets $t \geq c$).

Proof. The idea in the proof is to find a subset S of less than d memory cells that still allows us to answer $\geq d$ queries. This is impossible, since we can't generate d uniform outputs from less than d memory cells.

Let $p := 1/q^{1/4t}$. Include each memory bit in S with probability p , independently. By Theorem 2.8, $\mathbb{P}[|S| \geq cps] \leq 2^{-cps}$.

We are still able to answer a query if all of its memory bits fall in S . The probability that this happens is $p^t = 1/q^{1/4}$. We now claim that with probability $\geq 1/q^c$, we can still answer \sqrt{q} queries. Indeed, let B be the number of queries we cannot answer. We have $\mathbb{E}[|B|] \leq q(1 - q^{1/4})$. And so

$$\mathbb{P}[B \geq q(1 - 1/\sqrt{q})] \leq \frac{1 - q^{1/4}}{1 - \sqrt{q}} \leq 1 - q^c.$$

Thus, if the inequality $2^{-cps} \leq 1/q^c$ holds then there exists a set S of cps bits with which we can answer $\geq \sqrt{q} > d$ queries. Hence we reach a contradiction if

$$c \log q \leq cps < d.$$

Because $d > c \log q$ by assumption, and increasing s only make the problem easier, we reach a contradiction if $cps < d$, and the result follows. **QED**

Next we show a conceptually simple data structure which nearly matches the lower bound. For simplicity we focus on data structures which use space q^ϵ . Recall that for such space bounds the previous result does not give anything. We will show this is for good reasons, there are data structures where the time is constant. We will only get functions which are $c_\epsilon d$ -uniform, as opposed to d -uniform, but the proof techniques next and above generalize to other settings of parameters.

Theorem 12.2. There is a map $f : [q]^d \rightarrow [q]^q$ which is $c_\epsilon d$ -wise uniform and has a data structure with $w = \log q$ space $s = dq^\epsilon$ and time c_ϵ , for any ϵ and q which is a power of 2.

To give a sense of the parameters, let for example $q = d^{10}$.

Proof. We fill the memory with s evaluations of the input polynomial. Then we pick a random bipartite graph with s nodes on the left and q nodes on the right. Every node on the right side has degree g . We answer each query by summing the corresponding cells in s . Let $d' := d/g$. To show d' -wise uniformity it suffices to show that for any subset $R \subseteq [q]$ on the right-hand side of size d' , the sum of the corresponding memory cells is uniform in \mathbb{F}_q . For this in turn it suffices that R has a unique neighbor. And for that, finally, it suffices that R has a neighborhood of size greater than $\frac{g|R|}{2}$ (because if every element in the neighborhood of R has two neighbors in R then R has a neighborhood of size $< g|R|/2$).

Note here we are using that the neighborhood has size $\leq gd' = d$, and so the memory is d -wise uniform.

We pick the graph at random and show that it has the latter property with non-zero probability. We write $N(R)$ for the set of neighbors of R . We have:

$$\begin{aligned}
& \Pr [\exists R \subseteq [q], |R| \leq d', \text{ s.t. } |N(R)| \leq g|R|/2] \\
&= \Pr \left[\exists i \leq d', \exists R \subseteq [q], |R| = i, \text{ and } \exists T \subseteq [s], |T| \leq \frac{gi}{2} \text{ s.t. } N(R) \subseteq T \right] \\
&\leq \sum_{i=1}^{d'} \binom{q}{i} \binom{s}{gi/2} \left(\frac{gi/2}{s} \right)^{gi} \\
&\leq \sum_{i=1}^{d'} \left(\frac{eq}{i} \right)^i \left(\frac{es}{gi/2} \right)^{gi/2} \left(\frac{gi/2}{s} \right)^{gi} \\
&= \sum_{i=1}^{d'} \left(\frac{e^{1+g/2}q}{i} \right)^i \left(\frac{gi/2}{s} \right)^{gi/2} \\
&= \sum_{i=1}^{d'} \left[\underbrace{\frac{e^{1+g/2}q}{i} \left(\frac{gi/2}{s} \right)^{g/2}}_C \right]^i.
\end{aligned}$$

It suffices to have $C \leq 1/2$, so that the probability is strictly less than 1, because $\sum_{i=1}^k 1/2^i = 1 - 2^{-k}$. Recall that $gi \leq d$. Hence if $s = dq^\epsilon$ then $g = c_\epsilon$ suffices for large enough q . **QED**

12.1.1 Succinct data structures

Succinct data structures are those where the space is close to the minimum, n . Specifically, we let $s = n + r$ for some $r = o(n)$ called *redundancy*. Unsurprisingly, stronger bounds can be probed in this setting. But, surprisingly, again these stronger bounds were shown to be tight. Moreover, it was shown that improving the bounds would imply stronger circuit lower bounds.

To illustrate, consider the ECC problem $f : [2]^n \rightarrow [2]^m$ where f is an error-correcting code (with constant relative distance) and m is linear in n .

Theorem 12.3. Any data-structure for the ECC problem with $w = 1$ using space $n + r$ requires time $\geq cn/r$.

This is nearly matched by the following result.

Theorem 12.4. There is an ECC problem s.t. for any r it has a data structure with $w = 1$, space $n + r$, and time $c(n/r) \log^3 n$.

Moreover, it was shown that proving a time lower bound of $(n/r) \log^c n$ would imply new circuit lower bounds. The latter result refers to bounds on the number of wires in circuits with arbitrary gates. But the following connection with the standard circuit model is also known.

Theorem 12.5. Let $f : [2]^n \rightarrow [2]^{am}$ be a function computable by bounded fan-in circuits with bm wires and depth $b \log m$, for constants a, b . Then f has a data structure with space $n + o(n)$ and time $n^{o(1)}$.

Hence, proving n^ϵ time lower bounds for succinct data structures would give functions that cannot be computed by linear-size log-depth circuits, cf. 9.3.

12.1.2 Succincter: The trits problem

In this section we present a cute and fundamental data-structure problem with a shocking and counterintuitive solution. The trits problem is to compute $f : [3]^n \rightarrow ([2]^2)^n$ where on input n “trits” (i.e., ternary elements) $(t_1, t_2, \dots, t_n) \in [3]^n$ f outputs their representations using two bits per trit.

Example 12.1. For $n = 1$, we have $f(0) = 00, f(1) = 01, f(2) = 10$.

Note that the input ranges over 3^n elements, and so the minimum space of the data structure is $s = \lceil \log_2 3^n \rceil = \lceil n \log_2 3 \rceil \approx n \cdot 1.584 \dots$. This will be our benchmark for space. One can encode the input to f as before using bits without loss of generality, but the current choice simplifies the exposition.

Simple solutions:

- The simplest solution (cf. 12.2) to this problem is to use 2 bits per t_i . With such an encoding we can retrieve each $t_i \in [3]$ by reading just 2 bits (which is optimal). The space used is $s = 2n$ and we have linear redundancy.
- Another solution (cf. again 12.2) to this problem is what is called *arithmetic coding*: we think of the concatenated elements as forming a ternary number between 0 and $3^n - 1$, and we write down its binary representation. To retrieve t_i it seems we need to read all the input bits, but the space needed is optimal.
- For this and other problems, we can trade between these two extreme as follows. Group the t_i 's into blocks of t . Encode each block with arithmetic coding. The retrieval time will be ct bits and the needed space will be $(n/t)[t \log_2 3] \leq n \log_2 3 + n/t$ (assuming t divides n). This is block-wise arithmetic coding. It provides a *power* trade-off between retrieval time and redundancy. (Using number-theoretic results on logarithmic forms, one can show [210] that this last inequality is tight up to changing n/t into n/t^c .)

The shocking solution: An exponential (!) trade-off

We now present an *exponential* trade-off: retrieval time ct bits and redundancy $n/2^t + c$. In particular, if we set $t = c \log n$, we get retrieval time $c \log n$ and redundancy c . Moreover, the bits read are all consecutive, so with word size $w = \log n$ this can be implemented in constant time. To repeat, we can encode the trits with constant redundancy and retrieve each in constant time. This solution can also be made dynamic.

Theorem 12.6. The trits problem has a data structure with space $n \log_2 3 + n/2^t + c$ (i.e., redundancy $n/2^t + c$) and time ct , for any t and with word size $w = 1$. For word wise $w = \log n$ the time is constant.

Next we present the proof.

Definition 12.3 (Encoding and redundancy). An encoding of a set A into a set B is a one-to-one (a.k.a. injective) map $f : A \rightarrow B$. The *redundancy* of the encoding f is $\log_2 |B| - \log_2 |A|$.

The following lemma gives the building-block encoding we will use.

Lemma 12.1. For all sets X and Y , there is an integer b , a set K and an encoding

$$f : (X \times Y) \rightarrow ([2]^b \times K)$$

such that (1) f has redundancy $\leq c/\sqrt{|Y|}$, and (2) $x \in X$ can be recovered just by reading the b bits in $f(x, y)$.

Note that (1) says that $b + \log |K| - \log |X| - \log |Y| \leq c/\sqrt{|Y|}$. For (2) to hold we must have $b \geq \log |X|$. Combining this with the previous expression we obtain $\log |K| - \log |Y| \leq$

$c/\sqrt{|Y|}$. In particular we get that $|K| \leq 2^c \cdot |Y|$ (in fact it will be the case that $|K| \leq c \cdot \sqrt{|Y|}$, but the looser bound is sufficient).

The basic idea for proving the lemma is to break Y into $C \times K$ and then encode $X \times C$ by using b bits:

$$X \times Y \rightarrow X \times C \times K \rightarrow [2]^b \times K.$$

There is however a subtle point. If we insist on always having $|C|$ equal to, say, $\sqrt{|Y|}$ or some other quantity, then one can cook up sets that make us waste a lot (i.e., almost one bit) of space. The same of course happens in the more basic approach that just sets $Y = K$ and encodes all of X with b bits. The main idea will be to “reason backwards,” i.e., we will first pick b and then try to stuff as much as possible inside $[2]^b$. Still, our choice of b will make $|C|$ about $\sqrt{|Y|}$.

Proof. Pick any two sets X and Y , where $|Y| > 1$ without loss of generality. Define $b := \lceil \log_2 (|X| \cdot \sqrt{|Y|}) \rceil$, and let $B := [2]^b$. To simplify notation, define $d := 2^b/|X|$. Note $c\sqrt{|Y|} \leq d \leq c\sqrt{|Y|}$.

How much can we stuff into B ? For a set C of size $|C| = \lfloor |B|/|X| \rfloor$, we can encode elements from $X \times C$ in B . The redundancy of such an encoding can be bounded as follows:

$$\begin{aligned} & \log |B| - \log |X| - \log |C| \\ &= \log \frac{2^b}{|X|} - \log \lfloor \frac{2^b}{|X|} \rfloor \\ &= \log d - \log \lfloor d \rfloor \\ &\leq \log d - \log(d-1) \\ &= \log \left(1 + \frac{1}{d-1} \right) \\ &\leq \frac{c}{d-1} \\ &\leq \frac{c}{\sqrt{|Y|}-1} \\ &\leq \frac{c}{\sqrt{|Y|}}. \end{aligned}$$

To calculate the total redundancy, we still need to examine the encoding from Y to $C \times K$. Choose K of size $|K| = \lceil |Y|/|C| \rceil$, so that this encoding is possible. With a calculation similar to the previous one, we see that the redundancy is:

$$\begin{aligned}
& \log |C| + \log |K| - \log |Y| \\
&= \log \left\lceil \frac{|Y|}{|C|} \right\rceil - \log \frac{|Y|}{|C|} \\
&\leq \log \left(1 + \frac{|C|}{|Y|} \right) \\
&\leq c \frac{|C|}{|Y|} \\
&\leq c \frac{\lfloor \frac{2^b}{|X|} \rfloor}{|Y|} \\
&\leq c \frac{2^b}{|X| \cdot |Y|} \\
&\leq c 2^{(\log |X| \cdot \sqrt{|Y|})+1} / (|X| \cdot |Y|) \leq c \frac{\sqrt{|Y|}}{|Y|} = c \frac{1}{\sqrt{|Y|}}.
\end{aligned}$$

The total redundancy is then $c/\sqrt{|Y|}$, which gives (1).

For (2), it is clear from the construction that any $x \in X$ can be recovered from the element of B only. **QED**

Proof of Theorem 12.6. Break the ternary elements into blocks of size t : $(t'_1, t'_2, \dots, t'_{n/t}) \in T_1 \times T_2 \times \dots \times T_{n/t}$, where $T_i = [3]^t$ for all i . The encoding, illustrated in Figure 1, is constructed as follows, where we use f_L to refer to the encoding guaranteed by Lemma Lemma 12.1.

Compute $f_L(t'_1, t'_2) = (b_1, k_1) \in B_1 \times K_1$.

For $i = 2, \dots, n/t - 1$ compute $f_L(k_{i-1}, t'_{i+1}) := (b_i, k_i) \in B_i \times K_i$.

Encode $k_{n/t-1}$ in binary as $b_{n/t}$ using arithmetic coding.

The final encoding is $(b_1, b_2, \dots, b_{n/t})$. We now compute the redundancy and retrieval time.

Redundancy: From (1) in Lemma 12.1, the first $n/t - 1$ encodings have redundancy $c3^{-t/2} \leq 1/2^{ct}$. For the last (arithmetic) encoding, the redundancy is at most 1. So the total redundancy is at most $\left(\frac{n}{t} - 1\right) \cdot \frac{1}{2^{ct}} + 1 = \frac{n}{2^{ct}} + c$. One can visualize this as a “hybrid argument” transforming a product of blocks of ternary elements into a product of blocks of binary elements, one block at the time.

Retrieval Time: Say that we want to recover some t_j which is in block t'_i . To recover block t'_i , Lemma 12.1 guarantees that we only need to look at b_{i-1} and b_i . This is because k_{i-1} can be recovered by reading only b_i , and t'_i can be recovered by reading k_{i-1} and b_{i-1} . Thus to complete the proof it suffices to show that each b_i has length ct .

This is not completely obvious because one might have thought that the K_i become larger and larger, and so we apply the lemma to larger and larger inputs and the B_i get large too.

However, recall that each $|K_i| \leq c|T_i| = c3^t$ from the comment after the statement of Lemma 12.1. Hence, every time we apply the lemma on an input of size at most $s \leq 3^{ct}$. Since the lemma wastes little entropy (by (1) in Lemma 12.1), none of its outputs can be much larger than its input, and so $|B_i| = 2^{ct}$. **QED**

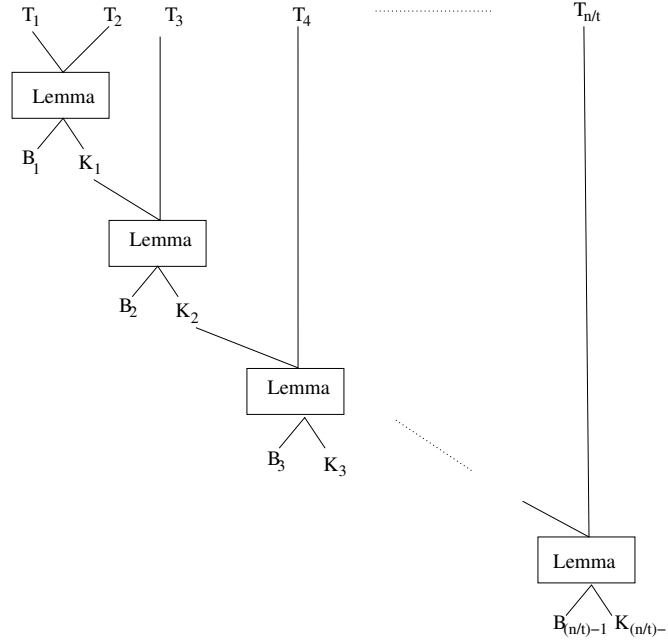


Figure 12.1: Succinct Encoding

12.2 Dynamic data structures

We now study dynamic data structures. As we mentioned, here the input is not fixed but can be modified by the queries.

Definition 12.4. Fix an error-correcting code $ECC : [2]^n \rightarrow [2]^m$ where $m \leq cn$ and $\Delta(ECC(x), ECC(y)) \geq c$ for any $x \neq y$ in $[2]^n$. Here $\Delta(u, v)$ is the relative distance, the fraction of bit positions where u and v differ.

The ECC problem asks to support operations, starting with the all-zero message:

$M(i, b)$ for $i \in \{1, 2, \dots, n\}$ and $b \in [2]$ which sets bit i of the message to b , and

$C(i)$ for $i \in \{1, 2, \dots, m\}$ which returns bit i of the codeword corresponding to the current message.

The time of a dynamic data structure is the maximum number of read/write operations in memory cells required to support an operation.

Theorem 12.7. The ECC problem requires time $t \geq c \log_w n \geq (c \log n) / \log \log n$ for cell size $w := \log n$.

One might wonder if stronger bounds can be shown for this problem. But in fact there exist codes for which the bounds are nearly tight.

Theorem 12.8. There exists codes for which the ECC problem can be solved in time $c \log^2 n$ with cell size $w = 1$.

Proof of Theorem 12.7. Pick $x \in [2]^n$ uniformly and $i \in \{1, 2, \dots, m\}$ uniformly, and consider the sequence of operations

$$M(1, x_1), M(2, x_2), \dots, M(n, x_n), C(i).$$

That is, we set the message to a uniform x one bit at a time, and then ask for a uniformly selected bit of the codeword $\text{ECC}(x)$, which we also denote by $C_x = C_x(1), C_x(2), \dots, C_x(n)$.

We divide the n operations $M(i, x_i)$ into consecutive blocks, called *epochs*. Epoch e consists of n/w^{3e} operations. Hence we can have at least $E := c \log_w n$ epochs, and we can assume that we have exactly this many epochs (by discarding some bits of n if necessary).

The geometrically decaying size of epochs is chosen so that the number of message bits set during an epoch e is much more than all the cells written by the data structure in future epochs.

A key idea of the proof is to see what happens when the cells written during a certain epoch are ignored, or reverted to their contents right before the epoch. Specifically, after the execution of the M operations, we can associate to each memory cell the last epoch during which this cell was written. Let $D^e(x)$ denote the memory cells of the data structure after the first n operations M , but with the change that the cells that were written last during epoch e are replaced with their contents right before epoch e . Define $C_x^e(i)$ to be the result of the data structure algorithm for $C(i)$ on $D^e(x)$, and $C_x^e = C_x^e(1), C_x^e(2), \dots, C_x^e(n)$.

Let $t(x, i, e)$ equal 1 if $C(i)$, executed after the first n operations M , reads a cell that was last written in epoch e , and 0 otherwise. We have

$$t \geq \max_{x,i} \sum_e t(x, i, e) \geq \mathbb{E}_{x,i} \sum_e t(x, i, e) = \sum_e \mathbb{E}_{x,i} t(x, i, e) \geq \sum_e \mathbb{E}_x \Delta(C_x, C_x^e), \quad (12.1)$$

where the last inequality holds because $C_x^e(i) \neq C_x(i)$ implies $t(x, i, e) \geq 1$.

We now claim that if $t \leq w$ then $\mathbb{E}_x \Delta(C_x, C_x^e) \geq c$ for every e . This concludes the proof.

In the remainder we justify the claim. Fix arbitrarily the bits of x set before Epoch e . For a uniform setting of the remaining bits of x , note that the message ranges over at least

$$2^{n/w^{3e}}$$

codewords. On the other hand, we claim that C_x^e ranges over much fewer strings. Indeed, the total number of cells written in all epochs after e is at most

$$t \sum_{i \geq e+1} n/w^{3i} \leq ctn/w^{3(e+1)}.$$

We can describe all these cells by writing down their indices and contents using $B := ctn/w^{3e+2}$ bits. Note that this information can depend on the operations performed during Epoch e , but the point is that it takes few possible values overall. Since the cells last changed during Epoch e are reverted to their contents before Epoch e , this information suffices to describe $D^e(x)$, and hence C_x^e . Therefore, C_x^e ranges over $\leq 2^B$ strings.

For each string in the range of C_x^e at most two codewords can have relative distance $\leq c$, for else you'd have two codewords at distance $\leq 2c$, violating the distance of the code.

Hence except with probability $2 \cdot 2^B / 2^{n/w^{3e}}$ over x , we have $\Delta(C_x, C_x^e) \geq c$. If $t_M \leq w$ then the first probability is ≤ 0.1 , and so $\mathbb{E}_x \Delta(C_x, C_x^e) \geq c$, proving the claim. **QED**

Exercise 12.4. Explain how to conclude the proof given the claim.

12.3 Notes

Theorem 12.1 is from [176]. It was rediscovered in [119].

Theorem 12.3: [63]. Efficient data structures for ECC, both static in Theorem 12.4 and dynamic in Theorem 12.8, as well as the connection to circuits, are from [214].

The breakthrough result on the trits problem is from [151]. After that a negative result was proved in [210], whose parameters were then matched by Theorem 12.6, proved in [151, 55]. Our exposition is from [205].

For dynamic data structures, the technique in the proof of Theorem 12.7 is from [60] and can be applied to many other natural problems, leading to tight results in several cases, see Exercise ???. It is not far from the state-of-the art in this area, which is $\log^{1+c} n$ [120].

Chapter 13

Pseudorandomness



Suppose I say to you that I've tossed coin 40 times and got this sequence of heads (0) and tails (1):

01..

You'd probably think this can't be true. But suppose instead I claim that I got

1000111110100010011110100101111101100100.

Maybe you would think this is possible then? But why do we feel this way? For a fair coin, the two strings have the same probability of 2^{-20} !

This example leads to a very interesting question: What is randomness? Of course, we've been using randomness all along since the first chapter. Still, let's step back and consider three possible answers:

1. **Classical:** Each string has the familiar probability. This viewpoint is useful in mathematics but the example above shows that it doesn't capture our intuitive notion of randomness.
2. **Intrinsic (or ontological):** A string is the less random the shorter description it has. The first string has the short program "Print 01 for 20 times," while the shortest program for the second seems to be "Print 1000111110100010011110100101111101100100."
3. **Behaviouristic:** Randomness is in the eyes of the beholder: A string R is random for f if f can't distinguish it from a truly random string. In other words, R fools f into thinking that R is random.

The last answer seems the most useful, and we now make it precise.

Definition 13.1. A distribution R over $[2]^n$ ϵ -fools (or is ϵ -pseudorandom for) a function $f : [2]^n \rightarrow [2]$ if $|\mathbb{E}[f(R)] - \mathbb{E}[f(U)]| \leq \epsilon$, where U is uniform in $[2]^n$. A function f ϵ -breaks (or tells, distinguishes) distributions D and E if $|\mathbb{E}[f(D)] - \mathbb{E}[f(E)]| \geq \epsilon$. If E is omitted it is assumed to be the uniform distribution.

We are naturally interested in distributions that are pseudorandom yet have very little entropy. As in Chapter 3, counting arguments show that very little entropy is needed for non-explicit distributions, about logarithmic in the number of tests to be fooled. But our ability to explicitly construct such distributions is limited by the grand challenge:

Claim 13.1. Let R be a distribution over $[2]^n$ with support of size $< 2^n/2$. Suppose that R $1/2$ -fools a set F of functions. Then the indicator function $g : [2]^n \rightarrow [2]$ of the support of R is not in F .

Proof. We have $\mathbb{E}[g(U)] < 1/2$ while $\mathbb{E}[g(R)] = 1$, so g is not $1/2$ -fooled and cannot be in F . **QED**

For example, if $F = \text{PCkt}$ and R can be sampled in P then $g \in \text{NP}$, and so $\text{NP} \not\subseteq \text{PCkt}$.

The above claim can be strengthened. In general, constructing such distributions can be thought of as a refined impossibility result that is closely related to correlation, recall Definition 10.3.

To simplify the following discussion, we introduce the notion of a pseudorandom generator which makes it easier to talk about the entropy of the distribution and its explicitness.

Definition 13.2. An algorithm G is a *pseudorandom generator*, abbreviated generator or PRG, that ϵ -fools a class F of functions (or a generator for F with error ϵ) with seed length s (a function of both n and ϵ) if on input n and ϵ , and a uniform seed U of length $s(n, \epsilon)$, G outputs a distribution on n bits that ϵ -fools any function in F on inputs of length n . The *stretch* is $n - s(n, \epsilon)$.

Note that we use n to denote the *output length* of G , because it is the *input length* for a test that's trying to tell G from random. Also, we typically have the output length of G much longer than the input length. For some applications, it suffices if G is computable in polynomial time in the output length, which can be exponential in the input length. We shall simply say that G is *explicit* in this case. However many generators we present below, especially those for restricted models, are explicit in a stronger sense: Given an input and an index to an output bit, that bit can be computed in P.

Exercise 13.1. Suppose there is $a > 0$ and an explicit generator with seed length $s(n) = a \log n$ that 0.1 -fools circuits of size n , for a constant a . Prove that $\text{P} = \text{BPP}$.

Note that to eliminate one parameter we set the size of the circuit test equal to the output length of G . Recall from Definition 2.4 that input gates are not counted towards size; the

circuit may simply ignore most of its input bits, which makes sense since very few input bits suffice, information-theoretically, to tell the output of G from uniform.

Given Claim 13.1, there are two main avenues for research, closely paralleling the development of earlier chapters. The first is proving unconditional results for restricted models, like AC. Actually, pseudorandomness being a more refined notion of impossibility, even very simple models like local functions are non-trivial, and results for them very useful. The second is proving reductions, that is linking the existence of PRGs to other conjectures. Interestingly, some of the techniques are general and apply to both settings.

13.1 Basic PRGs

In this section we present PRGs for several basic classes of tests. Besides being basic, these tests are the backbone of several other constructions, and somewhat surprisingly suffice to fool apparently stronger classes of tests.

13.1.1 Local tests

The simplest model to consider is perhaps that of local functions. A distribution over $[2]^n$ is *k-wise uniform* if every k bits are uniform in $[2]^k$ (equivalently, any k -local function is 0-fooled).

Exercise 13.2. Given an explicit 1-wise uniform generator with seed length $s(n) = 1$.

Theorem 13.1. There are explicit k -wise uniform generators with seed length $s = ck \log n$.

Proof. Wlog assume n is a power of 2 and let \mathbb{F} be the field of size n . More generally, the range of G will be \mathbb{F}^n , and the distribution of any k coordinates will be uniform over \mathbb{F}^k . View the input as coefficients a_i $i \in [k]$ of a polynomial p of degree $k - 1$. Define the i output element of G to be $p(i)$. Any k -tuple of field elements is uniform, for if two different polynomials give the same tuple then their difference is a non-zero polynomial of degree $k - 1$ with $\geq k$ roots, violating Lemma 2.3. (We can assume $k \leq n$ for else the theorem is trivial.)

QED

The bound on s is almost tight for small k . To see this, think of the support as $\{-1, 1\}^n$, and write down a $2^s \times n$ matrix where row x is $G(x)$. For even k , and for any $T \subseteq [n]$ of size $k/2$, consider the 2^s -long vector v_T obtained by multiplying together the columns indexed in T . Note that the v_T are orthonormal, hence independent, and so $2^s \geq \binom{n}{k/2}$, whence $s \geq ck \log(2n/k)$.

Powerlog-wise uniformity suffices to fool AC.

Theorem 13.2. Any $\log(m/\epsilon)^{cd}$ -wise uniform distribution over $[2]^n$ ϵ -fools AC of size m and depth d .

In particular, there are explicit generators that ϵ -fool such circuits with seed length $\log(m/\epsilon)^{cd}$. We give generators with this seed length below, via a different route.

13.1.2 Low-degree polynomials

Another natural model is that of low-degree polynomials. Chapter 6 and Chapter 10 give several applications, and we encounter more below in section 13.1.3.

Theorem 13.3. There are explicit generators that ϵ -fool degree-1 polynomials over \mathbb{F}_2 with seed length $s = c \log(n/\epsilon)$.

Exercise 13.3. Prove Theorem 13.3 using the construction in the proof of Lemma 6.6.

To fool polynomials of degree $d > 1$, we can take the xor of d independent copies of generators for degree 1. This is known to work for $d < \log n$, and is unknown beyond that.

Theorem 13.4. The sum of d generators that ϵ -fool degree-1 polynomials over \mathbb{F}_2 , on independent seeds, fools degree- d polynomials with error $\leq c\epsilon^{1/2^{d-1}}$.

Question 13.1. *Does this work for $d > \log n$?*

13.1.3 Expander graphs and combinatorial rectangles: Fooling AND of sets

In this subsection we construct a PRG to fool the And of (the indicator functions) of sets. We use “set” and “function” interchangeably in this section. This basic construction ties together many things we have seen and showcases techniques which allow to build even more powerful PRGs. Also, it suffices for the time-efficient simulation of BPP in PH, Item (2) in Theorem 6.3.

Theorem 13.5. There is a PRG G that ϵ -fools the product of t subsets of $[2]^m$ with seed length $m + c(\log t) \log(mt/\epsilon)$. In other words, for any functions $f_i : [2]^m \rightarrow [2]$ we have $|\mathbb{E}[\prod_i f_i(U_i)] - \mathbb{E}[\prod_i f_i(X_i)]| \leq \epsilon$ where $(X_1, X_2, \dots, X_t) = G(U)$.

Except for the extra $\log t$ factor, the seed length is good.

Exercise 13.4. Prove Item (2) in Theorem 6.3 assuming Theorem 13.5.

The fundamental case of $t = 2$ is known as *expander graphs*.

Exercise 13.5. [Where is the graph, and why is it expanding?] Let L and R be two disjoint sets with $M := 2^m$ nodes each, and define the graph on vertices $L \cup R$ and edges (x, y) from L to R for any output $(x, y) = G(z)$. Prove that any set $X \subseteq L$ of αM nodes has $\geq (1 - \epsilon/\alpha)M$ neighbors in R .

For expander graphs, explicit constructions with seed length $m + c \log 1/\epsilon$ are known. We give below a simpler construction with seed length $m + c \log(m/\epsilon)$. Expander graphs have many applications. A simple example is that the general case $t > 2$ is obtained from the $t = 2$ case via recursion.

Recursion

To fool $2t$ sets, first run the generator for 2 sets with error ϵ/c to get two seeds for generators for t sets with error ϵ/c . Then, run twice the generator for t sets on those seeds. Specifically, given $2t$ functions f_i , let $g_1 : [2]^{mt} \rightarrow [2]$ be the product of the first t , and g_2 the product of the last t . Let G_t be a generator for the product of t functions. We have:

$$|\mathbb{E}[g_1(U) \cdot g_2(U)] - \mathbb{E}[g_1(G(S_1)) \cdot g_2(G(S_2))]| \leq \epsilon/2.$$

To see this, define the “hybrid” distribution $H = g_1(G(S_1)) \cdot g_2(U)$, and note that the distance of $\mathbb{E}[H]$ from each of the expectations inside the absolute value is $\leq \epsilon/c$, and use the triangle inequality.

Now the key idea is that we can think of g_1 composed with G as another function h_1 , and similarly for g_2 . We can fool $h_1 \cdot h_2$ with the generator for two sets with error $\epsilon/2$, obtaining a generator for t sets with error $\epsilon/2 + \epsilon/2 = \epsilon$, as desired.

To analyze the seed length, denote it by $s(m, t, \epsilon)$ for parameters m , t , and ϵ . The definition above gives the recursion

$$s(m, 2t, \epsilon) \leq s(s(m, t, \epsilon/c), 2, \epsilon/c) \leq s(m, t, \epsilon/c) + c \log s(m, t, \epsilon/c) / \epsilon \leq s(m, t, \epsilon/c) + c \log(mt/\epsilon).$$

The second inequality is by the base $t = 2$ case, and the next is because seed mt always suffices, trivially. Iterating $\log_2 t$ times, we obtain seed length

$$\leq s(m, 2, \epsilon/t^c) + c \log(t) \log(mt/\epsilon)$$

which is as desired, using again the base case.

Expander graphs

The generator for the base case outputs $(U, U+D)$ where U is uniform and D is a distribution that ϵ -fools linear polynomials over \mathbb{F}_2 (!). By Theorem 13.3 the seed length is as desired. To analyze, it is natural to write the functions f_1 and f_2 in terms of polynomials. For slight convenience we think of the inputs in $\{-1, 1\}$ instead of $\{0, 1\}$, so that multiplication of input bits corresponds to xoring. In particular we will write $U \cdot D$ for $U + D$.

Exercise 13.6. For $\alpha \subseteq [n]$, we write x^α for $\prod_{i \in \alpha} x_i$. Let $f : \{-1, 1\}^n \rightarrow \mathbb{R}$ be a function.

(1) Show that f can be written as

$$f(x) = \sum_{\alpha} \hat{f}_{\alpha} \cdot x^{\alpha},$$

for some $\hat{f}_{\alpha} \in \mathbb{R}$. Guideline: First write $f(x) = \sum_{a \in \{-1, 1\}} f(a) I_a(x)$, where $I_a(x) = 1$ if $x = a$ and 0 otherwise.

(2) Show that $\hat{f}_{\emptyset} = \mathbb{E}[f(U)]$.

(3) Show that $\sum_{\alpha} \hat{f}_{\alpha}^2 = \mathbb{E}[f^2(U)]$.

Writing $f = f_1$ and $g = f_2$ we need to bound $|\mathbb{E}[f(U)f(U \cdot D)] - \mathbb{E}[f(U)]\mathbb{E}[g(U)]|$. Using Exercise 13.6, the second summand is $\hat{f}_\emptyset \hat{g}_\emptyset$. The first is

$$\mathbb{E}_{x \leftarrow U, D} \left[\sum_{\alpha, \beta} \hat{f}_\alpha \hat{g}_\beta x^\alpha (x \cdot D)^\beta \right].$$

Because $(x \cdot D)^\beta = x^\beta \cdot D^\beta$, the terms with $\alpha \neq \beta$ give 0. So we can rewrite it as

$$\mathbb{E}_D \left[\sum_{\alpha} \hat{f}_\alpha \hat{g}_\alpha D^\alpha \right].$$

Putting this together, we remove the $\alpha = \emptyset$ term, and our goal is to bound

$$\left| \mathbb{E}_D \left[\sum_{\alpha \neq \emptyset} \hat{f}_\alpha \hat{g}_\alpha D^\alpha \right] \right|.$$

This is at most

$$\begin{aligned} & \sum_{\alpha \neq \emptyset} |\hat{f}_\alpha| \cdot |\hat{g}_\alpha| \cdot |\mathbb{E}_D[D^\alpha]| \text{ (by the triangle inequality)} \\ & \leq \epsilon \sum_{\alpha} |\hat{f}_\alpha| \cdot |\hat{g}_\alpha| \text{ (by the assumption on } D) \\ & \leq \epsilon \sqrt{\sum_{\alpha} \hat{f}_\alpha^2} \cdot \sqrt{\sum_{\alpha} \hat{g}_\alpha^2} \text{ (by the inequality } \sum_i a_i b_i \leq (\sum_i a_i^2)(\sum_i b_i^2)) \\ & = \epsilon \sqrt{\mathbb{E}[f^2(U)]} \cdot \sqrt{\mathbb{E}[g^2(U)]} \text{ (by 13.6, Item (3))} \\ & \leq \epsilon \text{ (because the range of } f \text{ and } g \text{ is } [2]). \end{aligned}$$

13.2 PRGs from hard functions

In this section we present a general paradigm to construct PRGs from hard functions. We begin with a general claim showing that PRGs with non-trivial seed length $s(n) = n - 1$ are in fact equivalent to correlation bounds, recall Definition 10.3.

Claim 13.2. Let $f : [2]^n \rightarrow [2]$ be a function. We have:

(1) If $C : [2]^{n+1} \rightarrow [2]$ ϵ -breaks $G(x) := xf(x)$ then there is $b \in [2]$ s.t. $C'_b : [2]^n \rightarrow [2]$ defined as $C'_b(x) := C(xb) \oplus b$ has correlation $\mathbb{E}e[C'_b(x) \oplus f(x)] \geq \epsilon$ with f .

(2) Conversely, suppose $C : [2]^n \rightarrow [2]$ has ϵ -correlation with f . Then $C' : [2]^{n+1} \rightarrow [2]$ defined as $C'(x, b) := C(x) \oplus b$ ϵ -breaks $xf(x)$.

Proof of (1). Pick b uniformly and write

$$\mathbb{E}_{x,b}e[C'_b(x) \oplus f(x)] = \frac{1}{2}|\mathbb{E}_x C(xf(x)) - \mathbb{E}_x C(x\bar{f}(x))| = |\mathbb{E}_x C(xf(x)) - \mathbb{E}C(U)|.$$

So there is b s.t. the LHS is at least the RHS. This establishes the first claim. **QED**

Exercise 13.7. Prove (2) in Claim 13.2.

The contrapositive of (1) is that functions with small correlation immediately imply a 1-bit of stretch generator. Naturally, we'd like to increase the stretch. A natural idea is *repetition*: From a pseudorandom distribution D over $[2]^n$, we construct $D^k := D, D, \dots, D$ over $[2]^{k \cdot n}$.

Claim 13.3. If f ϵ -distinguishes D^k and E^k then a restriction of f ϵ/k -distinguishes D and E .

Proof. Via the “hybrid method,” a.k.a. the triangle inequality, cf. proof of Theorem 13.5. Define $H_i := D_0 D_1 \cdots D_{i-1} E_i E_{i+1} \cdots E_{k-1}$ over nk bits for $i \in [k]$, where each factor in the RHS is over n bits. Note that H_0 is E^k and H_k is D^k . Write

$$\begin{aligned} \epsilon &\leq |\mathbb{E}[f(H_0)] - \mathbb{E}[f(H_{k-1})]| \\ &= \left| \sum_{i \in [k]} \mathbb{E}[f(H_i)] - \mathbb{E}[f(H_{i+1})] \right| \\ &\leq \sum_{i \in [k]} |\mathbb{E}[f(H_i)] - \mathbb{E}[f(H_{i+1})]|. \end{aligned}$$

So one of the terms on the RHS is $\geq \epsilon/k$. The corresponding distributions H_i and H_{i+1} differ in only one factor. We can fix all others and the claim follows. **QED**

Note we went from ϵ to ϵ/k . This means the claim is only applicable when ϵ is fairly small. In general, this loss cannot be avoided:

Exercise 13.8. Exhibit a distribution D that is 0.1-pseudorandom (for say PCkt) but D^k is not even 0.9 pseudorandom, for suitable k . Now strengthen this to D of the form $xf(x)$, for some boolean function f .

However, repetition works for *resamplable* functions, like parity. These are functions h s.t. given any “correct” pair $(x, h(x))$ we can generate uniform correct pairs $(y, h(y))$, and similarly for “incorrect” pairs $(x, h(x) \oplus 1)$ – using the same distribution.

Definition 13.3. A function $h : [2]^n \rightarrow [2]$ is *resampled* by a distribution F on functions from $[2]^{n+1}$ to $[2]^{n+1}$ if for every $x \in [2]^n$ and $b \in [2]$, $F(x, h(x) \oplus b)$ outputs $(y, h(y) \oplus b)$ for uniform $y \in [2]^n$.

Claim 13.4. Suppose $h : [2]^n \rightarrow [2]$ is balanced (i.e., $\mathbb{P}[h(U) = 1] = 1/2$) and resampled by F . Let $D = (X, h(X))$. Suppose f ϵ -breaks D^k . Then $f(G, G, \dots, G)$ $\epsilon/2$ -breaks D , where each occurrence of G is either an occurrence of F or a fixed value.

Proof. Because h is balanced, we can sample U_{n+1} by first tossing a coin b , and then outputting $(X, h(X) \oplus b)$. Because f ϵ -breaks D , we can fix coins b_1, \dots, b_k s.t. the quantities

$$\begin{aligned} & \mathbb{E}[f((X_1, h(X_1)), (X_2, h(X_2)), \dots, (X_k, h(X_k)))] \\ & \mathbb{E}[f((X_1, h(X_1) \oplus b_1), (X_2, h(X_2) \oplus b_2), \dots, (X_k, h(X_k) \oplus b_k))] \end{aligned}$$

have distance $\geq \epsilon$.

The coordinates where $b_i = 0$ are the same. So we can fix those and obtain a restriction f' of f s.t. for some $j \leq k$ the quantities

$$\begin{aligned} & \mathbb{E}[f'((X_1, h(X_1)), (X_2, h(X_2)), \dots, (X_j, h(X_j)))] =: (I) \\ & \mathbb{E}[f'((X_1, \bar{h}(X_1)), (X_2, \bar{h}(X_2)), \dots, (X_j, \bar{h}(X_j)))] =: (II) \end{aligned}$$

have distance $\geq \epsilon$.

Now we use this to break D . As in the proof of Claim 13.2 it suffices to tell D from $\bar{D} := (X, \bar{h}(X))$. On input $z \in [2]^{n+1}$, we compute

$$f'(F(z), F(z), \dots, F(z)).$$

If $z = D$ then this is the same as (I), and if $z = \bar{D}$ this is the same as (II). **QED**

Claim 13.5. Parity is resamplable in AC.

Exercise 13.9. Prove this.

Combining the results in this section with the correlation of ACs and parity – Corollary 10.3 – we obtain a PRG with seed length $n - n/\log^{cd} n$ that fools ACs of size n^d and depth d on n bits. In the next section, leveraging the exponentially-small correlation bounds between ACs and parity, we will obtain a much shorter, logarithmic seed length for ACs.

However, for other classes of circuits like AC[2] such strong correlation bounds are not known. For these classes, the results in this section give the best-known explicit generator. For example, for AC[3] we can again use that parity has correlation $\leq 1/100$ with such circuits, as follows from Exercise 10.7, and obtain a generator stretching $n - n/\log^{cd} n$ bits to n . For AC[2] one can work with a different function and again obtain that stretch. Better stretch is not known.

Question 13.2. Give an explicit generator with seed length $0.9n$ for AC[2] circuits of size n^c and depth c on n bits.

13.2.1 Turning correlation bounds into stretch: Families of sets with small intersections

The repetition PRG outputs values of a hard function h on independent inputs. We now study a powerful technique which instead outputs values from *dependent* inputs. This gives a better trade-off between seed and output length. It is a derandomized analogue of Claim 13.3. Rather than picking independent inputs as in the repetition generator, we select them based on a collection of subsets of $[u]$, where u is the seed length.

Definition 13.4. Let $S = T_1, T_2, \dots, T_{|S|}$ be collection of subsets of $[u]$ of size ℓ . Then the *bounded-intersection generator*

$$\text{BIG}_S : [2]^u \rightarrow ([2]^\ell)^{|S|}$$

is defined as $\text{BIG}_S(x) := x_{T_1}, x_{T_2}, \dots, x_{T_{|S|}}$. Here x_{T_j} are the ℓ bits of x indexed by T_j .

For a distribution H on functions from $[2]^\ell \rightarrow [2]$ and a generator $G : [2]^u \rightarrow ([2]^\ell)^{|S|}$ we write $H \circ G(\sigma)$ for the result $H(x_1), H(x_2), \dots, H(x_{|S|})$ of applying H to the outputs of G , where $G(\sigma) = (x_1, x_2, \dots, x_{|S|})$ and the occurrences of H denote independent samples.

For example, if H is a uniform function then $H \circ \text{BIG}_S$ is uniform over $[2]^{|S|}$. The next key result shows that BIG preserves the indistinguishability of functions, similar to the repetition generator, as long as the sets in S have small intersections. The intersection size governs the locality (recall Definition 4.3), and hence the complexity, of the reduction.

Theorem 13.6. Let BIG and S be as in Definition 13.4. Furthermore, suppose $|T_i \cap T_j| \leq w$ for any $i \neq j$. Let V and W be two distributions on functions from $[2]^\ell$ to $[2]$.

Suppose f ϵ -tells the distributions $(\sigma, V \circ \text{BIG}_S(\sigma))$ and $(\sigma, W \circ \text{BIG}_S(\sigma))$, over $u + |S|$ bits. Then there are w -local functions g_i s.t. $f(g_1, g_2, \dots, g_{u+|S|})$ distinguishes $(X, V(X))$ from $(X, W(X))$ with advantage $\geq \epsilon/|S|$, where X is uniform in $[2]^\ell$.

Exercise 13.10. Derive Claim 13.3 from Theorem 13.6 for the special case $D = (X, V(X))$ and $E = (X, W(X))$.

Proof. Write $D = (\sigma, V \circ G_S(\sigma))$ and $E = (\sigma, W \circ G_S(\sigma))$. As in the proof of Claim 13.3, define hybrids

$$H_i := D_0 D_1 \cdots D_{i-1} E_i E_{i+1} \cdots E_{|\sigma|+|S|-1}$$

over $|\sigma| + |S|$ bits. Note that H_0 is E and $H_{|S|}$ is D . So there is i s.t. f distinguishes two adjacent hybrids H_i and H_{i+1} with advantage $\geq \epsilon/|S|$. We can fix the $u - \ell$ bits in the seed σ that are not in set T_i . Now every bit j in position $< i$ depends on $\leq w$ bits in T_i , and so can be computed by a distribution G_j on w -local functions.

The following distribution on circuits tells $(X, V(X))$ from $(X, W(X))$, again with advantage $\geq \epsilon/|S|$: On input (x, b) run f on

$$(G_0(x), G_1(x), \dots, G_{i-1}(x), b, G_{i+1}, G_{i+2}, \dots, G_{|\sigma|+|S|-1}).$$

We can fix the G_i to g_i and maintain the advantage. **QED**

To apply Theorem 13.6 we need a collection S with small intersections. We'd like to have as many sets as possible (that's the output length of the generator) which are as large as possible (that's the input length to the hard function) which are subsets of as small a set as possible (that's the seed length) and such that any two have as small intersection as possible (that's the overhead in the reduction). The probabilistic method shows that collections with great parameters exist. The following is a simple construction.

Lemma 13.1. [Sets with small intersections] There are explicit collections of q^d subsets of $[q^2]$ of size q such that any two sets intersect in $\leq d$ elements, for any q that is a power of 2 and $> d$.

Proof. View the universe $[u] = [q^2]$ as \mathbb{F}_q^2 . For a parameter d , the sets correspond to the graphs of polynomials p of degree $< d$. (I.e., the set $\{(x, p(x)) : x \in \mathbb{F}\}$.) The number of sets is q^d . The size of each set is $q = \sqrt{u}$. To bound the intersection of two sets, consider the corresponding polynomials and let p be their difference, which is not zero. Any element in the intersection of the sets corresponds to a zero of p . By Lemma 2.3, the intersection has size $\leq d$. **QED**

To illustrate parameters, we can have $m = q^d$ subsets of size $\ell = \sqrt{q}$ from a universe of size $u = \ell^2 = q$ with intersections at most d . For example, given m we can set $d = \log m$ and $q = \log^a m$, and the intersection size is only $\ell^{1/a}$ while the universe is only quadratic in the set size, i.e., $u = \ell^2$.

Corollary 13.1. There are generators $G : [2]^{\log^{c_d} n} \rightarrow [2]^n$ that $1/n$ -fool ACs of size n and depth d .

As in Exercise 13.1, in this corollary, to eliminate one parameter we set the size of the circuit equal to the output length of G . The same statement holds if the size is n^d instead of n .

Exercise 13.11. Prove Corollary 13.1. Explain how the parameters are set and which results you are combining.

The seed length in Corollary 13.1 is about the best we can do given current impossibility results, and recall once again from section §7.3 that stronger impossibility would imply major separations.

Still, one can ask if PRGs *could* be built *if* we had such stronger results. In particular, one would like to have seed length say $c \log n$ instead of $\log^c n$. This is the setting that allows for conclusions such as $P = BPP$, cf. Exercise 13.1. Lemma 13.1 doesn't give this, since the universe is always at least quadratic in the set size, but the following construction does.

Lemma 13.2. [Sets with small intersections, II] For any a and $n \geq c_a$ there is an explicit collection of n subsets of $[c_a \log n]$ of size $c_a \log n$ with pairwise intersection $\leq a \log n$.

Using this we can prove the following weaker version of Theorem 2.16.

Corollary 13.2. Suppose there is $\epsilon > 0$ and $f \in \mathbb{E}$ that on inputs of length n has correlation at most $2^{-\epsilon n}$ with circuits of size $2^{\epsilon n}$. Then $\text{P} = \text{BPP}$.

Exercise 13.12. Prove Corollary 13.2 assuming Lemma 13.2.

13.2.2 Turning hardness into correlation bounds

We can't expect to prove that correlation bounds under uniform are equivalent to impossibility or hardness results, as one can construct pathological functions which are easy to compute on, say, .75 fraction of the inputs, but impossible to compute on a .76 fraction. So instead our approach will be to *construct* functions which have small correlation under the uniform distribution.

A natural candidate for such a function, starting from a “mildly hard” function $f : [2]^n \rightarrow [2]$ is $f' : [2]^{nk} \rightarrow [2]$ defined as

$$f'(x_1, \dots, x_k) := \bigoplus_{i=1}^k f(x_i).$$

An *XOR Lemma* is a statement showing that if f has correlation $\leq \epsilon$ with a certain computational model (e.g., PCkt), then the correlation of f' with a related model decays *exponentially fast* with the number k of copies. There is a strong *information-theoretic* intuition why the XOR Lemma should work. Suppose each occurrence of $f(x_i)$ is a random bit $X_i \in [2]$ with $\mathbb{E}e[X_i] \leq \epsilon$ (cf. Definition 10.3). Then indeed $\mathbb{E}e[\sum X_i] = (\mathbb{E}e[X_1])^k \leq \epsilon^k$. Analogously, if f has correlation $\leq \epsilon$ with small circuits, then f' indeed has correlation $\leq \epsilon^k$ with small circuits *of the special product form* $C(x_1, \dots, x_k) := \bigoplus_{i=1}^k C_i(x_i)$. Intuitively, we can't do better than computing as in the special form. But is it true?

Exercise 13.13. Consider circuits C made of a single majority gate. Prove that the XOR lemma is false for C . Feel free to pick n even and define the value of Majority on inputs of weight $n/2$ to be 1, and recall $\binom{n}{n/2} \cdot \frac{\sqrt{n}}{2^n} \in [c, c]$.

One can extend this result to AC with a small number of majority gates.

Question 13.3. *Does the XOR lemma hold for AC with parity gates, or even constant-degree polynomials over \mathbb{F}_2 ?*

But for more powerful models, we can indeed prove the xor lemma, and the proof follows the information-theoretic intuition above. To connect to this intuition, we consider functions which may output a uniform bit on some inputs.

Definition 13.5. We say that a distribution on functions $F : [2]^n \rightarrow [2]$ is δ -*random* if there exists a subset $H \subseteq [2]^n$ with $|H| = 2\delta 2^n$ such that $F(x) = U_1$ (i.e. a coin flip) for $x \in H$ and $F(x)$ is deterministic (i.e., a fixed value) for $x \notin H$.

Thus, a δ -random function has a set of relative size 2δ on which it is information-theoretically unpredictable. To illustrate the XOR lemma, suppose that f is δ -random. Then f' will be almost a coin flip. Specifically, the probability that the output is not a coin flip is $(1 - 2\delta)^k$, the probability that no input falls into H . When some input falls into H , the output is a coin flip, and no circuit, efficient or not, can have non-zero correlation.

This intuition can be formalized via the *hardcore-set* lemma, which allows us to pass from computational hardness to information-theoretic hardness. Before stating the lemma we emphasize an important point:

The hardcore-set lemma is only known to hold for computational models which can compute majority. This is because the proof of correctness uses majority, as will be apparent in section §13.3. So to apply it, we have to start from an impossibility result for circuits that can compute majority. As discussed in Chapter 10, we essentially have no such result. In fact, in some restricted models, the xor lemma is false (cf. Exercise 13.13). So the results in this section are mostly conditional. Still, they allow us to spin a fascinating web of reductions between correlation and randomness, pointing to several challenges.

The following hardcore set lemma says that any function that has somewhat small correlation with small circuits admits a somewhat large hardcore set on which the function has very small correlation with small circuits. To illustrate parameters, note that a δ -random function has correlation $\leq 1 - 2\delta$ with any fixed function (or circuit) (we can extend Definition 10.3 to random functions by taking expectation over both the input and the random function). This is because when the input falls in the set H of density 2δ from Definition 13.5 then the correlation is zero. The hard-core set lemma shows that this is the only way that small correlation may arise: any function with small correlation with small circuits is in fact close to a δ -random function. We state the result in terms of distinguishing input-output pairs, as opposed to computing the function. This is equivalent by an argument similar to Claim 13.2 but is more convenient as it immediately allows us to talk about multiple inputs, as we also do in the next statement.

Lemma 13.3. Let $f : [2]^n \rightarrow [2]$ have correlation $\leq 1 - 2\delta$ with circuits of size s . Then there exists a $c\delta$ -random function $g : [2]^n \rightarrow [2]$ such that $X \cdot f(X)$ and $X \cdot g(X)$ are ϵ -indistinguishable for size $cs\epsilon^2\delta^2$, for any ϵ , where $X \equiv U_n$.

In particular, by Claim 13.3,

$$X_1 \cdots X_k \cdot f(X_1) \cdots f(X_k) \text{ and } X_1 \cdots X_k \cdot g(X_1) \cdots g(X_k)$$

are $k\epsilon$ -indistinguishable for size $cs\epsilon^2\delta^2$, where the X_i 's are uniform and independent.

We can now easily formalize the proof of the xor lemma.

Lemma 13.4. Suppose $f : [2]^n \rightarrow [2]$ is δ -hard for size s . Then $f' : [2]^{nk} \rightarrow [2]$ defined as $f'(x_1, \dots, x_k) := \bigoplus_{i=1}^k f(x_i)$ has correlation $\leq (1 - c\delta)^k + k/s^c$ with circuits of size $\delta^c s^c$.

For example, if $s = 2^{n^c}$ and $\delta = c$, we can take $k = cn$ and have hardness 2^{-cn} . However the function is on cn^c bits, so in terms of the input length n' , f' has hardness $2^{-n'^c}$.

Proof. We use Lemma 13.3 with $\epsilon := s^c$. From its conclusion it follows that

$$X_1 \cdots X_k \cdot \oplus_i f(X_i) \text{ and } X_1 \cdots X_k \cdot \oplus_i g(X_i)$$

are k/s^c -indistinguishable for size $\delta^c s^c$. Following the intuition above, the right-hand distribution is $(1 - c\delta)^k$ close to $X_1 \cdots X_k \cdot U_1$. Hence the left-hand distribution is $((1 - c\delta)^k + k/s^c)$ -indistinguishable from $X_1 \cdots X_k \cdot U_1$ and the result follows from Claim 13.2. **QED**

13.2.3 Derandomizing the XOR lemma

A drawback of the xor lemma is that the input length of the new function is $\geq kn$. This prevents us from obtaining correlation 2^{-cn} (as opposed to $2^{-c\sqrt{n}}$) which is important for the flagship conclusion $P = BPP$, cf. Corollary 13.2. To remedy this we shall use... PRGs! Rather than independently, we will pick the k inputs to f' using a generator. We need two properties from this PRG. First, to behave like repetition, we need BIG (Theorem 13.6). Also, we need to “hit” the hard-core set, for which we need HIT. We can get both properties by xor-ing the generators together. The generator is defined as

$$\text{BIG-HIT}(\sigma_1, \sigma_2) := \text{BIG}_S(\sigma_1) \oplus \text{HIT}(\sigma_2),$$

where HIT is a hitter, given next.

Lemma 13.5. For every ϵ and δ there exists an explicit generator $\text{HIT} : [2]^{2n} \rightarrow ([2]^n)^s$ with $s = 1/\epsilon\delta$ s.t. for every set $H \subseteq [2]^n$ of size ϵ , $\mathbb{P}_\sigma[\text{HIT}(\sigma)_i \notin H \text{ for every } i] \leq \delta$.

Proof. Pairwise independence. Consider the field \mathbb{F}_2 . The seed σ specifies $a, b \in \mathbb{F}$ and we output $b, a + b, 2a + b, \dots$. Let X_i be the indicator variable of $\text{HIT}(\sigma)_i \in H$. The X_i are pairwise independent. Their expectation is ϵs . Hence the probability to bound is $\leq \mathbb{P}[|\sum X_i - \epsilon s| \geq \epsilon s]$. Squaring both sides of the inequalities and doing calculations gives the result. **QED**

Exercise 13.14. Do the calculations.

Using this, we can boost correlation $(1 - 2^{-cn})$ to correlation $\leq 2^{-cn}$. We give an example for an interesting setting of parameters.

Lemma 13.6. Suppose E has a function $f : [2]^* \rightarrow [2]$ that on inputs of length n has correlation $(1 - 2^{-cn})$ with circuits of size 2^{cn} . Then E has a function $f' : [2]^* \rightarrow [2]$ that has correlation $\leq 2^{-cn}$ with circuits of size 2^{cn} .

Note the conclusion implies $P = BPP$ by Corollary 13.2.

Proof. Let $\epsilon := 2^{-cn}$ and $\delta := 2^{-cn}$. Define $f' : [2]^{cn} \rightarrow [2]$ as $f'(\sigma) := \bigoplus_{i=1}^s f(x_i)$ where $BIG-HIT(\sigma) = (x_1, x_2, \dots, x_s)$, where $s = 1/\delta\epsilon$ and the set system for BIG is from Lemma 13.2.

We use Lemma 13.3 with $\epsilon := s^c$. Let g the corresponding δ -random function. From Theorem 13.6 it follows that

$$\sigma, f \circ G \text{ and } \sigma, g \circ G$$

are ϵ^c -indistinguishable for size $1/\epsilon^c$. In particular this holds if we take parities, so

$$\sigma, \bigoplus_{i=1}^k f(X_i) \text{ and } \sigma, \bigoplus_{i=1}^k g(X_i)$$

are no more distinguishable, where $(X_1, \dots, X_k) = G(\sigma)$. By the hitting property of HIT, Lemma 13.5, the chance of not hitting the hardcore set is $\leq \delta$, and we conclude as in the proof of Lemma 13.4. **QED**

Exercise 13.15. Actually, this proof doesn't quite follow from Theorem 13.6 as stated. Explain why and why it isn't an issue.

13.2.4 Encoding the whole truth-table

The results in the previous section give us functions with small correlation starting from functions on n bits with with correlation $(1 - 2^{-cn})$, but not quite from any impossibility results, which only gives correlation $\leq (1 - 2/2^{-n}) < 1$.

Exercise 13.16. Explain where the previous proofs break down.

To start from worst-case hardness we need to encode the entire truth table of the function. We give a simple code that suffices for our results.

Theorem 13.7. Suppose there is $f \in E$ that on inputs of length n cannot be computed by circuits of size $s(n)$. Then there is $f' \in E$ that has correlation $(1 - 1/n^c)$ with circuits of size $n^c s(cn)$.

Proof. Let $q = n^{10}$ and $d = n^5$ and view the truth-table of f as specifying coefficients over \mathbb{F}_q for a polynomial in $\ell := n/\log n$ variables with the degree in each variables being $\leq d$. Note the number of monomials is $\geq d^\ell \geq 2^n$, larger than the truth-table of f .

We can identify f with the corresponding polynomial p_f . Via interpolation, we can define p_f so that evaluating f can be reduced to evaluating p_f .

The new function f' is constructed in two steps. First, we consider inputs over \mathbb{F}_q^ℓ . Note the length of such inputs is $\leq \ell q \leq cn$ bits, as desired. This gives a non-boolean function. To make the function boolean, we output bit i of p_f , where i is part of the new input. That is,

$$f'(x_1, \dots, x_\ell, i) := p_f(x_1, \dots, x_\ell)_i$$

where $x_i \in \mathbb{F}_q$ and $i \in [\log q]$.

We'd like to show that if there's a small circuit C computing f' on a $(1 - 1/n^c)$ fraction of inputs then there's another small circuit computing f everywhere. Let $C(x) := C(x, 1) \cdots C(x, \log q)$. First note that the fraction α of $x \in \mathbb{F}_q^\ell$ such that $C(x) \neq p_f(x)$ is $\leq 1/n^c \leq c/d\ell$. Because if it's larger, every such x contributes at least one input (x, i) where C disagrees with f' , contradicting the assumption.

Using C we give a distribution C' on circuits which computes p_f w.h.p. on every given input y . Pick a uniform line going through y , and run C on this line for $d\ell$ points. That is, pick uniform $s \in \mathbb{F}_q^\ell$ and run $C(y + 1s), C(y + 2s), \dots, C(y + d\ell s)$.

Because each evaluation point is uniform, and $d\ell\alpha \leq c$, with prob. $> 1/2$ the evaluations of C will be correct, and equal $p_f(y + 1s), p_f(y + 2s), \dots, p_f(y + d\ell s)$.

Note that for fixed y and s , $p_f(y + ts)$ is a univariate polynomial q in t of degree $\leq d\ell$. We can compute the coefficients of q from its evaluations at $d\ell$ points. (It's a linear system, with a unique solution by Lemma 2.3 because the degree of q is $\leq \ell \cdot d < q$.)

We can then output $q(0) = p_f(y)$.

Finally, we can repeat this cn times and output the most likely value. On every input x this errs w.p. $< 2^{-n}$. Hence we can fix the random choices and obtain a fixed circuit that succeeds on every x . **QED**

Exercise 13.17. "Put it all together" and prove Theorem 2.16.

13.2.5 Monotone amplification within NP

To increase the hardness of functions in NP we cannot use XOR since NP is not known to be closed under complement. We will use a combination of many things in this chapter – including the (unconditional) generator for AC in Corollary 13.1 – to establish the following.

Theorem 13.8. If NP has a balanced function that has correlation $\leq 1/10$ with circuits of size 2^{n^c} , then NP also has a balanced function with correlation $\leq 2^{-n^c}$ with circuits of size 2^{n^c} .

Several optimizations have been devised, see the Notes. Still, we don't enjoy the same range as for E:

Question 13.4. Prove Lemma 13.6 for NP instead of E, even starting from hardness $\delta \geq c$.

13.2.6 Proof of Theorem 13.8

Rather than XOR, to amplify we use the Tribes function, a monotone read-once DNF.

Definition 13.6. The Tribes function on k bits is:

$$\text{Tribes}(x_1, \dots, x_k) := (x_1 \wedge \dots \wedge x_b) \vee (x_{b+1} \wedge \dots \wedge x_{2b}) \vee \dots \vee (x_{k-b+1} \wedge \dots \wedge x_k)$$

where there are k/b clauses each of size b , and b is the largest integer such that $(1 - 2^{-b})^{k/b} \geq 1/2$. Note that this makes $b \leq c \log k$.

The property of xor that we used is that if one bit is uniform, then the output is uniform. We use an analogous property for tribes, that if several bits are uniform, then the output is close to uniform.

Lemma 13.7. Let N_p be a noise vector where each is 1 independently with probability p . Then $\mathbb{E}_{x, N_p} e[\text{Tribes}(x) \oplus \text{Tribes}(x \oplus N_p)] \leq 1/k^{c_p}$.

We shall take k exponentially large. The resulting function is still in NP as we can use non-determinism to pick a clause. We use the generator $\text{BIG-AC}(\sigma) = (x_1, x_2, \dots, x_s)$, which is like BIG-HIT except that HIT is replaced with the generator in Corollary 13.1, for circuits of size $(k2^n)^c$. Note its seed length is n^c for $k \leq 2^{n^c}$.

Define $f' : [2]^{2^n} \rightarrow [2]$ as $f'(\sigma) := \text{Tribes} \circ (f(x_1), \dots, f(x_s))$ where $\text{BIG-AC}(\sigma) = (x_1, x_2, \dots, x_s)$, and the set system for BIG is from Theorem 3.1. Following the proof of Lemma 13.6, use Lemma 13.3 with $\epsilon := s^c$. Let g the corresponding δ -random function. From Theorem 13.6 it follows that

$$\sigma, f \circ \text{BIG-AC} \text{ and } \sigma, g \circ \text{BIG-AC}$$

are ϵ^c -indistinguishable for size $1/\epsilon^c$. In particular this holds if we take Tribes of the output, i.e. What remains to show is that

$$\sigma, \text{Tribes} \circ g \circ \text{BIG-AC}$$

is close to uniform. That is, we have to show that with high probability over σ , just over the choice of g , the value $\text{Tribes} \circ g \circ \text{BIG-AC}$ is close to a uniform bit.

It suffices to bound

$$\mathbb{E}_\sigma |\mathbb{E}_g e[\text{Tribes} \circ g \circ \text{BIG-AC}(\sigma)]|.$$

Here the inner expectation is over the random choices in all the s evaluations of g . Up to a power, this is

$$\leq \mathbb{E}_\sigma \mathbb{E}_g^2 e[\text{Tribes} \circ g \circ \text{BIG-AC}(\sigma)] = \mathbb{E}_{\sigma, g, g'} e[\text{Tribes} \circ g \circ \text{BIG-AC}(\sigma) \oplus \text{Tribes} \circ g' \circ \text{BIG-AC}(\sigma)].$$

Now the critical step is that $\text{Tribes} \circ g$ is computable by a distribution on AC of size $(s2^n)^c$ and depth c . Note that the circuit computes g in brute-force, but the dependence on s is good. Because BIG-AC fools such circuits with error 2^{-n^c} , the latter expectation equals

$$\mathbb{E}_{(x_1, \dots, x_s), g, g'} e[\text{Tribes} \circ g \circ (x_1, \dots, x_s) \oplus \text{Tribes} \circ g' \circ (x_1, \dots, x_s)] = \mathbb{E}_{x, N_p} e[\text{Tribes}(x) \oplus \text{Tribes}(x \oplus N_p)],$$

for $p = c$. We conclude by Lemma 13.7.

13.3 Hardcore distributions

In this section we put the hardcore set Lemma 13.3 in context, and prove it. Recall that the lemma establishes that for any hard function there is a distribution *which is uniform over a large set*, wrt which circuits have small correlation.

What if we just want *any* distributions? It turns out that then impossibility results are actually equivalent to correlation bounds! Consider a set F of functions mapping $[2]^n$ to $[n]$, for example, F could consist of circuits of a certain size. Also, let h be a target function, e.g., a function we aim to show is hard.

Corollary 13.3. Suppose for every distribution D on $[2]^n$ there is $f \in F$ s.t. $\mathbb{E}_{x \leftarrow D} [f(x) + h(x)] \geq \epsilon$. Then there exist cn/ϵ^2 functions $f_i \in F$ s.t.

$$h = \text{Majority}(f_1, f_2, \dots, f_{cn/\epsilon^2}).$$

In particular, for models that are able to compute majority, such as PCkt, we get that $h \notin \text{PCkt}$ iff there is a distribution D over $[2]^n$ s.t. any $f \in \text{PCkt}$ has correlation $\leq 1/n^a$ for any a . In other words, *superpower correlation bounds for some distribution are necessary and sufficient for superpower impossibility*.

We can give a converse of Corollary 13.3, and in fact even obtain correlation under any distribution, including uniform distribution.

Claim 13.6. Suppose $h = \text{Majority}(f_1, f_2, \dots, f_t)$. Let D be a distribution on inputs. Then there is i s.t. $\mathbb{E}[(f_i + h)(D)] \geq 1/t$.

Exercise 13.18. Prove this. Feel free to assume t is odd for simplicity.

Note that if we “only” had Corollary 13.3 when D is uniform, we could have skipped all the amplification results in the previous section, and constructed PRGs much more directly. However, again, in general we can’t guarantee that. In fact, one can construct functions that are very easy over the uniform distribution, say because they are almost always one, but still are hard to compute, say because there is a small set of inputs that makes the function hard.

Still, the techniques used to prove Corollary 13.3 are the same used to prove the hardcore set Lemma 13.3, which, recall, gives correlation bounds for the uniform distribution *over some large set*, and that in turn be used to obtain correlation bounds over the uniform distribution over the entire input space. So we now develop machinery to understand and prove Corollary 13.3. A useful viewpoint, here and elsewhere, is the *equivalence between having randomness in the input and having it in the model*:

Corollary 13.4. There is a distribution over F s.t. $\mathbb{E}_F [f(x) + h(x)] \geq \alpha$ for every x iff for every distribution D over $[2]^n$ there is $f \in F$ s.t. $\mathbb{E}_{x \leftarrow D} [f(x) + h(x)] \geq \alpha$.

Exercise 13.19. Prove the “only if” direction.

Corollary 13.4 is a special case of the min-max theorem from game theory, a.k.a. linear-programming duality, etc, stated next.

Theorem 13.9. Let X and Y be sets, $p : X \times Y \rightarrow \mathbb{R}$, and $\alpha \in \mathbb{R}$. Then either there is a distribution D_X on X s.t. $\mathbb{E}_{D_X} p(D_X, y) \geq \alpha$ for every $y \in Y$, or there is a distribution D_Y on Y s.t. $\mathbb{E}_{D_Y} p(x, D_Y) \leq \alpha$ for every $x \in X$.

To see the correspondence, we can let $X := [2]^n$ be the set of inputs, and $Y := F$ be the set of functions, and $p(x, f) = e[f(x) = h(x)]$. If for every distribution D over $[2]^n$ there is $f \in F$ that computes h well over D , the first condition in Theorem 13.9 does not hold. So the second holds.

Corollary 13.3 follows from Corollary 13.4 and tail bounds (Theorem 2.8), similarly to the proof of the error-reduction Theorem 2.9 for BPTIME.

Proof of 13.3. Use Theorem 13.9 to get a distribution on F . The majority of cn/ϵ^2 samples from F has error $< 2^{-n}$ by Theorem 2.8. By a union bound we can fix the samples to the f_i . **QED**

13.3.1 Proof of the hardcore-set Lemma 13.3

At the high-level, this is just min-max and concentration of measure, just like before. However, the proof is slightly more involved than one might anticipate. We break it up in the two claims. In the first we obtain hardness w.r.t. a “smooth” distribution D : $D(x) \leq d/N$ for every x . For example, D could be uniform over a set of size N/d (then $D(x)$ is either 0 or d/N). In the second we obtain a set from a smooth distribution. The straightforward combination of the claims yields the lemma.

Claim 13.7. Suppose $f : [2]^n \rightarrow [2]$ is $1/d$ -hard for circuits of size s . Then there is a distribution D on $[2]^n$ s.t. $D(x) \leq d/N$ for every x , and every circuit C of size $s \cdot (\epsilon/\log d)^c$ has $\mathbb{E}_{x \leftarrow D} e[C(x) + f(x)] \leq \epsilon$.

Proof. We use the min-max Theorem 13.9 where one set consists of sets S of N/d inputs, and the other consists of circuits of size $\leq s$, and $f(S, C) = \mathbb{E}_{x \in S} e[C(x) + f(x)]$.

Suppose there is a distribution D over sets of N/d inputs such that for every circuit we have $\mathbb{E}_D [f(D, C)] = \mathbb{E}_{S \leftarrow D} \mathbb{E}_{x \in S} e[C(x) + f(x)] \leq \epsilon$. Let D be the induced distribution over x and note that $D(y) \leq d/N$ for every y , and we’re done.

Otherwise by the min-max Theorem 13.9 there is a distribution C on circuits of size s s.t. for any set S of size N/d we have $\mathbb{E}_{S \leftarrow D} \mathbb{E}_C e[C(x) + f(x)] \geq \epsilon$. Let S be the set of inputs on which the inner expectation is $\leq \epsilon/2$. Note $|S| < (1 - \epsilon/2)N/d$, for else the set S' consisting of the N/d elements where the inner expectation is smallest would have only $(N/d)\epsilon/2$ elements outside of S , yielding expectation $< \epsilon/2 + \epsilon/2 = \epsilon$, contradiction. For every $x \notin S$, picking $\log(d)/\epsilon^c$ samples from C and taking majority gives error probability $\leq \epsilon/2d$, using Theorem 2.8 as in the proof of Theorem 2.9. The prob. of not computing correctly a uniform $x \in [N]$ is at most the prob. that $x \in S$ plus the prob. that the samples of C give the wrong value: $(1 - \epsilon/2)/d + \epsilon/2d = 1/d$. This contradicts the hardness of f .

QED

When using the following claim for a hard function h , we can let F be the set of functions of the type $e^{(h(x)+C(x))}$ where C is a small circuit. In this way $|\mathbb{E}_D[f(D)]|$ is the correlation of h and C w.r.t. D .

Claim 13.8. Let D be a distribution over $[N]$ s.t. $D(x) \leq d/N$. Let F be a set of $\leq ce^{c\epsilon^2 N/d^2}$ functions $f : [N] \rightarrow \{-1, 1\}$. Suppose for every $f \in F$ we have $\mathbb{E}_D[f(D)] \leq \epsilon$.

Then there is a set $S \subseteq [N]$ of size $|S| \geq cN/d$ s.t. for every $f \in F$ we have $\mathbb{E}_{x \in S}[f(x)] \leq c\epsilon$.

Even this second step is not immediate, due to the fact that the set S is constructed probabilistically and so its size – which is the normalization in the correlation – is not fixed. So we'll first prove concentration around a quantity related to D only, then connect it to $|S|$.

Proof. Construct S by placing each $x \in [N]$ in S independently with prob. $D(x)N/d \in [0, 1]$. Consider $X := \sum_{x \in [N]} S(x)f(x)$, where S is the indicator of set S . The variables $S(x)f(x)$ are independent and have range $[-1, 1]$. Also, $\mathbb{E}[X] = (N/d)\mathbb{E}_D[f(x)]$, and so $|\mathbb{E}[X]| \leq c\epsilon N/d$. By tail bounds, Exercise 2.16:

$$\mathbb{P}_S[|\sum_{x \in [N]} S(x)f(x)| \geq c\epsilon N/d] \leq 2e^{-c\epsilon^2 N/d^2}.$$

Also, $\mathbb{E}[\sum_{x \in [N]} S(x)] = N/d$. And so again by tail bounds the probability that $|S| \leq cN/d$ is, say, $\leq e^{-cN/d^2}$.

By a union bound, there exists S of size $\geq cN/d$ s.t. for every $f \in F$ we have

$$|\sum_{x \in [N]} S(x)f(x)| \leq c\epsilon N/d.$$

Now it's the moment to connect to $|S|$. Dividing both sides by $|S|$ we have

$$|\mathbb{E}_{x \in S}[f(x)]| \leq c\epsilon(N/d)/|S| \leq c\epsilon,$$

as desired. **QED**

Problem 13.1. Give a “direct” construction of PRGs sufficient to prove $P = BPP$ from a δ -hard function h in E . Guideline: Use BIG-HIT to generate an $n \times n$ matrix of inputs, evaluate h on every input, and then XOR the rows. Start with $\delta = c$. How small can you make δ and still have this construction?

13.4 Notes

For more on unconditional pseudorandom generators see [87]. For a broader view of pseudorandomness, with an emphasis on connections between various objects, see [194]. For Fourier analysis, see [147].

For expander graphs see [93]. They have many equivalent presentations, for example in terms of eigenvalues. My presentation is in terms of the mixing lemma, see Section 2.4 in [93]. In my definition I allow for repeated edges. Different notions of explicitness are also natural. In my definition one can output an edge given an index. More stringently, one can ask, given a node and an index to an incident edge, to compute the corresponding neighbor. The construction I presented immediately gives the more stringent explicitness as well.

k -wise uniform distributions were studied before complexity theory, cf. [155]. The complexity viewpoint is from [48, 12].

Generators for degree-1 polynomials originate in [137], with alternative constructions in [13]. The idea of xoring generators for degree-1 polynomials to fool higher-degree polynomials is from [36]. It was studied further in [126, 208], with the latter paper proving Theorem 13.4.

Regarding 13.2: A construction computable in time n^c first appeared in [145], where 13.2 also appears. Alternative constructions computable with small space or with alternations appeared respectively in [113] and [204]. Still, all these constructions use resources at least exponential in the seed length, while for several applications such as section ?? one needs power in the seed length. This stronger explicitness is obtained in [79] building on an idea presented in [82] of using error-correcting codes. Specifically, one can use the polynomial code from 2.21 in combination with 13.2 for logarithmic-scale collections (for which the former notion of explicitness is now acceptable).

The XOR lemma was reportedly announced in talks associated with the work [222], cf. [71]. Hardness amplification within NP was first studied in [146] which established correlation about $1/\sqrt{n}$. Exponentially small correlation was achieved in [88], with optimizations in [127, 74]. Our exposition follows [88].

Corollary 13.4 and the connection to min-max is from [224]. Hardcore sets were introduced in [98]. They were optimized and shown to be connected to boosting techniques in machine learning in [112] and subsequent works.

Other proofs of Theorem 2.16 don't use the derandomized xor lemma, or only use it from constant hardness, and instead rely on results in [98] (see e.g. the original proof [102]) or [187] (see e.g. [194] or [16]).

Problem 13.1 is similar to a construction in [187], except I use XOR instead of extractors, cf. Remark 15 in [187].

For more on hitters, see Appendix C in [67].

13.4.1 Myth creation: Polylogarithmic independence fools AC (Theorem 13.2)

In 1991 [143] constructed a pseudorandom generator for AC (a.k.a. alternating circuits or AC0 circuits), vastly improving the parameters of the pioneering work [7]. This is one of my favorite papers ever. (Mini myth creation episode: A large fraction of papers cite [145] for this result, possibly even the majority. This issue of credit is indeed complicated, since the 1988 conference version of [145] claims ownership for this AC result, and cites an unpublished manuscript with the same title as [143], but with both authors. One can only guess that the

authors decided that the AC result should only be attributed to Nisan.)

Nisan's distribution, and even the earlier one in [7], is polylog-wise uniform, that is, any polylog bits are uniform. (The polylog depends on the parameters of the circuit to be fooled in a standard way which is ignored here.) In fact, these results apply to a natural class of polylog-wise distributions: If you pick a uniform sparse linear transformation, the output distribution will be polylog-wise uniform, and Nisan's proof shows that it fools AC.

However, the proof does not show that *every* polylog-wise distribution fools AC. Later, Linial and Nisan [125] conjectured that polylog-wise uniformity suffices to fool AC, which would generalize both [7] and [143].

This problem was somewhat notorious but there was no progress until the paper by Bazzi [29], 15+ years after the conjecture was posed, which proves it for the special case of DNFs.

Bazzi's paper is quite hard to read, and the journal version is also long – 60 pages. Consequently it was hard to find referees, both for the conference and the journal version. Things must have gotten somewhat desperate, because when it finally was my turn to be asked to review the journal version it was deemed appropriate to extract a commitment from me before I could see the submission, something that has never occurred to me for any other paper. My back-and-forth with the author during the refereeing process was then abruptly stopped by, I suspect, the circulation of Razborov's follow up [161] which dramatically simplifies the presentation, especially with an idea by Wigderson. It was then clear that the results were correct and the paper could be accepted, even though I never claimed to understand Bazzi's proof for the non-monotone case.

The message in the papers [29] and [161] was loud and clear: You can make progress with just a little duality. From Razborov's paper:

By linear duality, this conjecture is an approximation problem of precisely the kind considered in [LMN93, BRS91, ABFR94]. Therefore, it is quite remarkable that the only noticeable progress in this direction was achieved only last year by Bazzi [Baz07].

At this point it was clear that the general case of AC might not be that hard. Shortly after Razborov's paper, Braverman [38] indeed proved this, albeit with a quadratic rather than linear dependence on the depth of the circuit. This dependence was later improved.

As usual we can look at citation count:

[38] 143

[28] 91

But more interestingly the literature is full of citations like:

A breakthrough result by Braverman [No mention of Bazzi or Razborov]

My definition of breakthrough result is roughly that of progress on a problem such that many people have thought about it but have been stuck for a long time. This applies to [28].

Approximate number of years gap:

[125]-[28]: XXXXXXXXXXXXXXXXXXXX

[28]-[161]: XX

[161]-[38]: X

I also think that if a problem was open even for depth 2, then going from 1 to 2 tends to be more fundamental than going from 2 to d . One can think of situations where this wouldn't be the case, for example if the depth-2 case was known for a while, and people were really stuck and couldn't do even depth 3, and that turned out to require a completely different approach. This isn't the case here.

Consider the following example. Tonight a breakthrough lower bound for depth-3 Majority circuits comes out. Then in a year this result is extended to any constant depth with additional but related techniques. Which result, if any, is the breakthrough?

13.5 Problems

Problem 13.2. Let $\text{ACSize}(d, s)$ be the functions computable by explicit ACs of size s and depth d . Prove that $\text{BP} \cdot \text{ACSize}(d, s) \subseteq \text{ACSize}(d + c, 2^{\log^{c_d} s})$.

Chapter 14

Communication complexity

This chapter deals with Communication Complexity, which is the study of the amount of information that needs to be exchanged among two or more parties (or players) which are interested in reaching a common computational goal.

14.1 Two parties

We start with the model in which there are only 2 parties, A and B . They have the following properties:

- They collaborate, and
- each party has unlimited computing power.

Their task is to compute a predefined function of two inputs

$$f : X \times Y \rightarrow [2]$$

where A only knows $x \in X$, and B only knows $y \in Y$. The parties A and B engage in a communication protocol and exchange bits. At each step, the protocol specifies whose turn is to speak, or if the protocol is over. This is a function of the bits exchanged so far. If a party is to speak, the protocol specifies which bit is sent, and this is a function of both the bits exchanged so far and the input to the party who is to speak. If the protocol is over, the last bit exchanged is the output. We say that the protocol uses d bits if for every input A and B exchange $\leq d$ bits. The bits exchanged are called *transcript*.

We can visualize a protocol via a binary tree (14.1). Each node is labeled with a party and a function from that party's input to $\{0, 1\}$, which specifies which children to go to.

14.1.1 The communication complexity of equality

Consider the function Equality : $[2]^n \times [2]^n \rightarrow [2]$, Equality(x, y) = 1 \Leftrightarrow $x = y$. Trivially, Equality can be computed with communication $n + 1$: A sends her input to B ; B then




Figure 14.1: Protocol tree

communicates the value of Equality. The same trivial upper bound holds for any function $f :: [2]^n \times [2]^n \rightarrow [2]$. We now prove the following lower bound.

Theorem 14.1. Any protocol for equality must exchange at least n bits.

Before proving this theorem, we cover some properties of protocols.

Definition 14.1. A rectangle in $X \times Y$ is a subset $R \subseteq X \times Y$ such that $R = A \times B$ for some $A \subseteq X$ and $B \subseteq Y$.

An equivalent definition is given by the following proposition

Proposition 14.1. $R \subseteq X \times Y$ is a rectangle iff $(x, y) \in R$ and $(x', y') \in R \Rightarrow (x, y'), (x', y) \in R$.

The connection between rectangles and protocols is the following.

Lemma 14.1. Let P be a protocol that uses d bits, let $t \in [2]^d$ be a transcript. The set of inputs that induce transcript t is a rectangle.

Proof. Let $A \subseteq X \times Y$ be the set of inputs that induce communication t . Suppose that $(x, y), (x', y') \in A$, we want to show that $(x, y') \in A$ (similarly for (x', y)). We prove by induction on i that the i -th bit exchanged by P on input (x, y') is t_i . Of course this means that the protocol exchanges t on input (x, y') and so $(x, y') \in A$ as desired.

For $i = 1$, the bit sent by A only depends on x , but we know $P(x, y)$ exchanges t_1 , so we are done.

For general i , suppose it is A 's turn to speak. The bit she sends is a function of x and the communication so far. By induction hypothesis the communication so far is t_1, \dots, t_{i-1} . So A cannot distinguish between (x, y) and (x, y') and will send t_i as next bit.

If it is B 's turn to speak, we reason in the same way replacing (x, y') with (x', y') . **QED**

Corollary 14.1. Suppose $f : X \times Y \rightarrow [2]$ is computable by a d -bit protocol, then there is a partition of $X \times Y$ in 2^d rectangles, where each rectangle is f -monochromatic: all inputs in the rectangle give the same value of f .

1	0	0	0
0	1	0	0
0	0	1	0
0	0	0	1

1	0		
	1	0	0
0		1	
	0	0	1

Figure 14.2: Two ways to partition equality in monochromatic rectangles.

Proof. For each transcript t , consider $R_t :=$ the set of inputs that induce t . R_t is a rectangle by the previous lemma. It is obviously a partition and f -monochromatic. **QED**

Figure 14.2 shows two ways to partition equality in monochromatic rectangles. We can now prove the lower bound for equality.

Proof of Theorem 14.1. Assume we can partition $X \times Y$ in equality-monochromatic rectangles. Consider the 2^n inputs (e, e) where $e \in \{0, 1\}^n$. Observe that no equality-monochromatic rectangle can contain both (e, e) and (b, b) if $e \neq b$, for else (e, b) is in the rectangle, but since $e \neq b$ this cannot be equality-monochromatic.

Since the rectangles must cover all of the 2^n inputs (e, e) , we need $\geq 2^n$ rectangles which implies that any protocol must use at least n bits of communication. **QED**

14.1.2 The power of randomness

We can define a randomized protocol as a distribution on protocols, cf. section §??. A function has randomized communication k with error ϵ if there is a randomized protocol that on every input computes it correctly w.p. $\geq 1 - \epsilon$. The equality function demonstrates the power of randomness in communication:

Theorem 14.2. Equality has randomized protocols with error ϵ and communication $c \log 1/\epsilon$, for any $\epsilon \leq 1/2$.

Exercise 14.1. Prove this.

14.1.3 Public vs. private coins

14.1.4 Disjointness

The *disjointness function* $\text{Disj} : [2]^n \times [2]^n \rightarrow [2]$ is defined as $\text{Disj}(x, y) = \bigvee_{i \in [n]} x_i \wedge y_i$. It asks to determine if x and y , viewed as subsets of $[n]$, (do not) intersect. This function is of central importance pretty much for the same reason that 3Sat is: Its simple structure makes it excellent for reductions, as we shall see in section 14.1.7.

Theorem 14.3. [107, 160] The randomized communication complexity of Disj with error ϵ is $\geq c_\epsilon n$.

The hard distribution D is defined as follows for $n = 4m - 1$. First pick a uniform partition of $[n]$ into $(P, Q, \{i\})$ where P (and Q) is a uniform set of size $2m - 1$. Now let X (resp., Y) be a uniform subset of $P \cup \{i\}$ (resp. $Q \cup \{i\}$) of size m . In particular, the intersection is either empty or a singleton. Note that the distribution is not product; it is known that the communication is $\leq c\sqrt{n}$ on product distributions.

14.1.5 Greater than

Another well-studied function is Greater-Than, where the parties wish to determine if $x > y$ as integers.

Theorem 14.4. [144] The randomized communication complexity of greater-than is $\leq c \log n$.

Proof. We sketch the clever protocol. We perform binary search to find the most significant bit where x and y differ. Each comparison during this binary search corresponds to an equality problem, which as we saw has small randomized communication complexity (Theorem 14.2).

The naive way to implement this search is to set the error to $\leq c/\log n$ in Theorem 14.2. But this won't give overall communication $c \log n$.

Instead, we set the error to constant, and perform *binary search with noisy comparisons*. A random-walk-with-backtrack algorithm [57] shows that $c \log n$ comparisons suffice, leading to the result. The idea is to start each recursive call with a check that the target element is contained in the current interval, and if not backtrack. **QED**

The above bound is tight.

Theorem 14.5. [212] The randomized communication complexity of greater-than is $\geq c \log n$.

14.1.6 Application to TMs

One-tape TMs have efficient randomized communication protocols. This is essentially the same as the crossing-sequence argument we saw in Chapter 3.

Theorem 14.6. For a function $f : [2]^n \times [2]^n \rightarrow [2]$ consider the padded function $p_f : [2]^{3n} \rightarrow [2]$ defined as $p_f(x0^n y) = f(x, y)$. If p_f is computable by an s -state TM in time t then f has randomized protocols with communication $c(\log s)t/n$ and error $\leq 1/2$.

Proof. For $i \in [n]$, define the protocol P_i as follows: A is in charge of the first $n + i$ cells (which include x); B is in charge of last $n + (n - i)$ cells (which include y). They simulate the TM in turn, communicating $\log s + c$ bits whenever the TM crosses the boundary of the $(n + i)$ -th cell. These bits represent the state of the machine or a special symbol denoting

that the computation is over with final state s , from which the value of the function can be determined. The parties carry this simulation for up to $(t/n)/(c \log s)$ crossings. If the TM hasn't stopped they stop and output, say, 0.


The distribution on protocols is P_I where I is uniform in $[n]$. **QED**

We will soon exhibit functions which require linear randomized communication, recovering the quadratic impossibility results for TMs from Chapter 3. In fact, we will show stronger results.

14.1.7 Application to streaming

14.2 Number-on-forehead

There are various ways in which we can generalize the 2-party model of communication complexity to $k > 2$ parties. The obvious generalization is to let k players compute a k -argument function $f(x_1, \dots, x_k)$ where the i -th party only knows the i -th argument x_i . This model is known as “number-in-hand” and useful in some scenarios, but we will focus on a different, fascinating model which has an unexpected variety of applications: the “Number on the Forehead” model [45]. Here, again $f(x_1, \dots, x_k)$ is a Boolean function whose input is k arguments, and there are k parties. The twist is that the i -th party knows all inputs except x_i , which we can imagine being placed on his forehead. Communication is broadcast.



The grand challenge here is to give an explicit function $f : \overbrace{[2]^n \times \dots \times [2]^n}^k \rightarrow [2]$ that cannot be computed with $k := 2 \log n$ parties exchanging k bits. This would have many applications, one of which is described next.

14.2.1 An application to ACC

Functions computable by small ACC (recall section §10.5) have low communication complexity:

Theorem 14.7. AC[d] on n bits of size n^d and depth d have equivalent protocols with $\log^{cd} n$ parties communicating $\log^{cd} n$ bits, for any partition of the input bits.

Proof. By Lemma 10.5 it suffices to prove it for depth-2 circuits consisting of a symmetric gate on s and gates of fan-in t , where s and t are $\leq \log^{cd} n$. Fix an arbitrary partition of the

input in $t + 1$ sets x_1, \dots, x_{t+1} . All that the players need to compute is the number of And gates that evaluate to 1. Consider any And gate. Since it depends on at most t variables, it does not depend on the bits in one of the sets, say x_j . Then the j -th party can compute this And without communication. So let us partition the And gates among the parties so that each party can compute the gates assigned to them without communication. Each party evaluates all the And gates assigned to them privately and broadcasts the number $\leq s$ of these gates that evaluate to 1. This takes a total of $ct \log s$ bits. **QED**

14.2.2 Generalized inner product is hard

In this section we prove an impossibility result for computing the generalized inner product function $GIP : ([2]^n)^k \rightarrow [2]$:

$$GIP(x_1, \dots, x_k) := \sum_{i=1}^n \bigwedge_{j=1}^k (x_j)_i \pmod{2}.$$

In fact, we shall bound even the correlation $\text{Cor}(GIP, d\text{-bit } k\text{-party})$ between GIP and k -party protocols exchanging d bits, defined as the maximum of $|\mathbb{E}_x e[GIP(x) + f(x)]|$ for any protocol f with corresponding parameters. Here we use $e(z) := (-1)^z$.

Theorem 14.8. [22] $\text{Cor}(GIP, d\text{-bit } k\text{-party}) \leq 2^d \cdot 2^{-\Omega(n/4^k)}$.

To prove the theorem we associate to any function a quantity $R(f) \in \mathbb{R}$ enjoying the following two lemmas:

Lemma 14.2. $\text{Cor}(f, d\text{-bit}) \leq 2^d \cdot R(f)^{1/2^k}$, for any $f : X_1 \times \dots \times X_k \rightarrow [2]$.

Lemma 14.3. $R(GIP) \leq 2^{-\Omega(n/2^k)}$.

The combination of these two facts proves Theorem 14.8.

Intuition for $R(f)$: Think of $k = 2$; we saw that any 2-party d -bit protocol partitions the inputs in 2^d f -monochromatic rectangles. How about we check how well f can be so partitioned? Instead of picking an arbitrary rectangle, let us pick one in which each side has length 2, and see how balanced the function is there. If a “good” partition exists, with somewhat high probability our little rectangle should fall in a monochromatic rectangle, and we should always get the same values of f . Otherwise, we should get mixed values of f .

Specifically, for $k = 2$,

$$R(f) := \mathbb{E}_{\substack{x_1^0, x_2^0 \\ x_1^1, x_2^1}} [f(x_1^0, x_2^0) + f(x_1^0, x_2^1) + f(x_1^1, x_2^0) + f(x_1^1, x_2^1)] \in \mathbb{R}.$$

In general, for any k :

$$R(f) := \mathbb{E}_{\substack{x_1^0, \dots, x_k^0 \\ x_1^1, \dots, x_k^1}} \left[\sum_{\varepsilon_1, \dots, \varepsilon_k \in [2]} f(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k}) \right] \in \mathbb{R}.$$

14.2.3 Proof of Lemma 14.2

We prove this theorem via a sequence of claims.

Definition 14.2. A function $g_i : X_1 \times \dots \times X_k \rightarrow [2]$ is a *cylinder* in the i -th dimension if $\forall (x_1, \dots, x_k)$ and x'_i we have $g_i(x_1, \dots, x_{i-1}, x_i, x_{i+1}, \dots, x_k) = g_i(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_k)$. A set $S \subseteq X_1 \times \dots \times X_k$ is a *cylinder intersection* if \exists cylinders g_1, \dots, g_k such that $S = \{x : \prod g_i(x) = 1\}$.

Recall we saw that a 2-party protocol partitions the input in monochromatic rectangles. The following extension of this fact to k parties is via cylinder intersections.

Claim 14.1. Any d -bit k -party protocol for $f : \overbrace{[2]^n \times \dots \times [2]^n}^k \rightarrow [2]$ partitions the inputs in 2^d f -monochromatic cylinder intersections.

Proof. Fix a transcript t , and consider the set A_t of inputs yielding that transcript. We claim that A_t is a cylinder intersection. To see this, consider the cylinder functions $g_i(x) = 1 \Leftrightarrow$ “From the point of view of the i -th party, x could yield transcript t ” $\Leftrightarrow \exists x'_i$ such that $P(x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_k)$ yields transcript t .

Obviously if x is in A_t then $g_i(x) = 1$ for all i .

To see the converse, take some input $x = (x_1, \dots, x_k)$ such that $g_i(x) = 1$ for all i . This means that $\exists (x'_1, \dots, x'_k)$ such that

- (x'_1, x_2, \dots, x_k) yields t ;
- (x_1, x'_2, \dots, x_k) yields t ;
-
- (x_1, x_2, \dots, x'_k) yields t .

We must show that x yields t as well, i.e. $x \in A_t$. This is argued by induction on the bits in t , using the same “copy and paste” argument that was used for $k = 2$. **QED**

Using the notion of cylinder intersections we can now relate an arbitrary protocol to a special class of protocols p^* . Each protocol p^* can be written as $p^*(x) = \sum g_i(x) \text{ mod } 2$, where g_i is a cylinder in i -th dimension. This corresponds to each party sending just one bit independently of the others, and the output of the protocol being the XOR of the bits. Note the communication parameter is not present anymore. We write Cor^* for the corresponding correlation, where k is given by the context.

Claim 14.2. $\text{Cor}(f, d\text{-bit}) \leq 2^d \cdot \text{Cor}^*(f)$.

Proof. We use a general trick to turn products $\overbrace{\prod_i g_i(x) = 1}^{\text{cylinder intersection}}$ into sums $\overbrace{\sum g_i(x) \text{ mod } 2}^{p^*}$. Fix any d -bit protocol, let $\{x : \prod_i g_i^1(x) = 1\}, \dots, \{x : \prod_i g_i^D(x) = 1\}$ be the corresponding

$D := 2^d$ f -monochromatic cylinder intersections (by the previous claim). Observe that for a fixed x ,

$$\mathbb{E}_{y_1, \dots, y_k \in \{-1, 1\}} \left[(y_1)^{1+g_1(x)} \cdot (y_2)^{1+g_2(x)} \cdot \dots \cdot (y_k)^{1+g_k(x)} \right] = \begin{cases} 1 & \text{if } \exists i : g_i(x) = 0 \\ 0 & \text{if } \forall i : g_i(x) = 1. \end{cases}$$

Therefore,

$$e(p(x)) = \sum_{i=1}^D r(i) \mathbb{E}_{y_1, \dots, y_k \in \{-1, 1\}} \left[(y_1)^{1+g_1^i(x)} \cdot (y_2)^{1+g_2^i(x)} \cdot \dots \cdot (y_k)^{1+g_k^i(x)} \right]$$

where $r(i) \in \{-1, 1\}$ is the value of the protocol on the i -th cylinder intersection. Note that for any x exactly one expectation will be 1, the one corresponding to the cylinder intersection where x lands. So we have:

$$\begin{aligned} & \mathbb{E}e[f(x) + p(x)] \\ &= \mathbb{E}_x[e(f(x)) \cdot e(p(x))] \\ &= \mathbb{E}_x \left[e(f(x)) \cdot \sum_{i=1}^D r(i) \mathbb{E}_{y_1, \dots, y_k \in \{-1, 1\}} \left[(y_1)^{1+g_1^i(x)} \cdot (y_2)^{1+g_2^i(x)} \cdot \dots \cdot (y_k)^{1+g_k^i(x)} \right] \right] \\ &= \sum_{i=1}^D \mathbb{E}_{x, y_1, \dots, y_k \in \{-1, 1\}} \left[e(f(x)) \cdot r(i) \cdot (y_1)^{1+g_1^i(x)} \cdot (y_2)^{1+g_2^i(x)} \cdot \dots \cdot (y_k)^{1+g_k^i(x)} \right] \\ &\leq D \cdot \mathbb{E}_{x, y_1, \dots, y_k \in \{-1, 1\}} \left[e(f(x)) \cdot r(i) \cdot (y_1)^{1+g_1^{i^*}(x)} \cdot (y_2)^{1+g_2^{i^*}(x)} \cdot \dots \cdot (y_k)^{1+g_k^{i^*}(x)} \right], \end{aligned}$$

where i^* is the value of i that gives the largest summand. Now fix y_1, \dots, y_k to maximize the expectation, and let $J \subseteq \{1, \dots, k\}$ be the indices corresponding to $y_j = -1$, i.e., $j \in J \Rightarrow y_j = -1$. The last expression above is

$$\begin{aligned} & D \cdot \mathbb{E}_x \left[e(f(x)) \cdot \prod_{j \in J} (-1)^{1+g_j^{i^*}(x)} \right] \\ &= D \cdot \mathbb{E}_x \left[e(f(x) + \sum_{j \in J} (1 + g_j^{i^*}(x))) \right] \\ &\leq D \cdot \text{Cor}^*(f). \end{aligned}$$

QED

Claim 14.3. $\mathbb{E}e_x[g(x)] \leq R(g)^{1/2^k}$ for every function $g := X_1 \times \dots \times X_k \rightarrow [2]$.

Proof. Recall that for every random variable X : $E[X^2] \geq E[X]^2$. This holds because $0 \leq E[(X - E[X])^2] = E[X^2] - E[X]^2$. Also recall that if X, X' are independent then $E[X \cdot X'] = E[X] \cdot E[X']$.

We proceed with the “squaring trick.”

$$\begin{aligned} \mathbb{E}_{x_1, \dots, x_k} e[g(x_1, \dots, x_k)]^2 &= \mathbb{E}_{x_1, \dots, x_{k-1}} [\mathbb{E}_{x_k} e[g(x_1, \dots, x_k)]]^2 \leq \mathbb{E}_{x_1, \dots, x_{k-1}} [\mathbb{E}_{x_k} e[g(x_1, \dots, x_k)]^2] \\ &= \mathbb{E}_{x_1, \dots, x_{k-1}} [\mathbb{E}_{x_k^0, x_k^1} e[g(x_1, \dots, x_{k-1}, x_k^0) + g(x_1, \dots, x_{k-1}, x_k^1)]]]. \end{aligned}$$

The lemma follows by repeating this k times. **QED**

Claim 14.4. For every function $f : X_1 \times \dots \times X_k \rightarrow [2]$, and every protocol* p^* ,

$$R(f \oplus p^*) = R(f),$$

where $f \oplus p^*$ simply is the function whose output is the XOR of f and p^* .

Proof. Suppose $p^*(x) = g_1(x) + \dots + g_k(x)$, where g_i is a cylinder in the i -th dimension. We show $\forall f, R(f \oplus g_k) = R(f)$; the same reasoning works for the other coordinates. Note for every x ,

$$\begin{aligned} & \sum_{\varepsilon_1, \dots, \varepsilon_k \in [2]} (f(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k}) + g_k(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k})) \\ &= \sum_{\varepsilon_1, \dots, \varepsilon_k} f(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k}) + \sum_{\varepsilon_1, \dots, \varepsilon_k} g_k(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k}) \\ &= \sum_{\varepsilon_1, \dots, \varepsilon_k} f(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k}) + \sum_{\varepsilon_1, \dots, \varepsilon_k} g_k(x_1^{\varepsilon_1}, \dots, x_k^0) \\ &= \sum_{\varepsilon_1, \dots, \varepsilon_k} f(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k}) + 2 \sum_{\varepsilon_1, \dots, \varepsilon_{k-1}} g_k(x_1^{\varepsilon_1}, \dots, x_k^0) \\ &= \sum_{\varepsilon_1, \dots, \varepsilon_k} f(x_1^{\varepsilon_1}, \dots, x_k^{\varepsilon_k}) \pmod{2}, \end{aligned}$$

where the second equality holds because g_k does not depend on x_k . **QED**

The straightforward combination of the claims in this section proves Lemma 14.2.

14.2.4 Proof of Lemma 14.3

We have:

$$\begin{aligned} R(\text{GIP}) &= \mathbb{E} e^{\sum_{\varepsilon_1, \dots, \varepsilon_k \in \{0,1\}} \sum_i \prod_j (x_j^{\varepsilon_j})_i} = \mathbb{E} \prod_i e^{\sum_{\varepsilon_1, \dots, \varepsilon_k} \prod_j (x_j^{\varepsilon_j})_i} \\ &= \mathbb{E} e^{\sum_{\varepsilon_1, \dots, \varepsilon_k} \prod_j (x_j^{\varepsilon_j})_1} = R\left(\bigwedge_k\right)^n, \end{aligned}$$

using in the last equality the fact that any two independent random variables X, Y satisfy $E[X \cdot Y] = E[X] \cdot E[Y]$, and where \bigwedge_k is the AND function on k bits.

To save in notation let us replace $(x_1^0)_1, \dots, (x_k^0)_1$ with (y_1^0, \dots, y_k^0) , where $(y_i^0) \in [2]$; and similarly for $(x_1^1)_1, \dots, (x_k^1)_1$. So we have:

$$R(\text{GIP}) = \mathbb{E} e^{\sum_{\varepsilon_1, \dots, \varepsilon_k \in \{0,1\}} \prod_j y_j^{\varepsilon_j}}.$$

Suppose that $y_1^0 \neq y_1^1, \dots, y_k^0 \neq y_k^1$; then there exists exactly one choice of $\varepsilon_1, \dots, \varepsilon_k$ making $\prod_j y_j^{\varepsilon_j} = 1$, and consequently

$$e^{\left(\sum_{\varepsilon_1, \dots, \varepsilon_k} \prod_j y_j^{\varepsilon_j}\right)} = e(1) = -1.$$

We have $y_1^0 \neq y_1^1, \dots, y_k^0 \neq y_k^1$ with probability 2^{-k} . Therefore:

$$R(\text{GIP}) = \mathbb{E} \left[\sum \prod_j y_j^{\varepsilon_j} \right]^n \leq (-1 \cdot 2^{-k} + 1 \cdot (1 - 2^{-k}))^n = (1 - 2^{-k+1})^n \leq e^{-cn/2^k}.$$

14.3 Efficient protocols with logarithmically many players

The impossibility results in the previous sections are effective when the number of players is $k \leq c \log n$, but useless when $k \geq \log n$. We now show that this is for a good reason: there are efficient protocols for large k . For generalized inner product this is unsurprising, since the function is almost always 0 for large k . But this is not clear if we replace, say, And with Majority. In fact, surprisingly there is a general protocol that works for many such “combined” functions.

Theorem 14.9. [19] Let $k \geq \log n + 2$. There is a k -party protocol with communication $c \log^2 n$ s.t. given a $k \times n$ matrix M (player j sees all M except row j) computes (y_0, y_1, \dots, y_n) where y_i is the number of columns in M with weight i .

In particular, any symmetric function of (the outputs of symmetric functions of the columns) has protocols with the same efficiency.

Proof. Each player j communicates the number $a_j(i)$ of columns that they see having weight i .

From this they can compute

$$b_i := \sum_j a_j(i)$$

for $i \in [t]$.

We claim that the b_i uniquely specify the y_i , which concludes the proof.

To verify the claim, note we have for $i \in [t]$:

$$(k - i)y_i + (i + 1)y_{i+1} = b_i.$$

Assume towards a contradiction that there are y_i and y'_i for $i \in [t + 1]$ that satisfy these equations, and are non-negative and have the same sum, n , and for some $i \in [t + 1]$ we have $y_i \neq y'_{i+1}$.

Let $d_i := y_i - y'_i$. Subtracting the equations above we get for $i \in [t]$

$$(k - i)d_i + (i + 1)d_{i+1} = 0.$$

One can verify by induction that the above implies

$$d_i = (-1)^i \binom{k}{i} d_0.$$

We know that $d_0 \neq 0$ (for else all the d_i would be 0) and in fact $d_0 \geq 1$ since it is an integer. Also note that $y_i + y'_i \geq |y_i - y'_i| = |d_i|$.

Hence we obtain the following contradiction

$$2n = \sum_{i=0}^k y_i + y'_i \geq |d_i| \geq \sum_{i=0}^k \binom{k}{i} = 2^k > 2n.$$

QED

This striking result has been generalized in multiple ways.

14.3.1 The power of randomness

A non-explicit linear separation is in [32]. An explicit, power separation is in [111]. A candidate for an explicit linear separation the problem of deciding if $xyz = 1_G$ for a group G . With randomness, this can be solved with constant communication, for any group.

Exercise 14.2. Show this.

Without randomness, there are groups where this requires $c \log \log |G| \geq c \log n$ communication, see [215]. An open question in this area is resolving whether bounds like $\geq c \log |G|$ hold for some group.

14.3.2 Pointer chasing

We consider the basic problem of following a path in a directed graph. We have k layers, with edges going from nodes in layer i to nodes in layer $i + 1$ only. The last layer does not have edges but labels in $[2]$ for each node. The goal is to compute the label at the node reached from a start node in Layer 1. (Equivalently, instead of labels we can allow for $k + 1$ layers with the last layer consisting of two nodes only, and the task is outputting the node reached.)

It is convenient to work with a version of this problem where we output m labels, and where the size of each layer grows by a factor b . Note for $m = 1$ this is a boolean function. We shall show that the communication is at least b .

Formally, the input to the pointer-chasing function $G_k^{m,b}$ is a layered graph as above, where Layer i has mb^i nodes, and each node has outdegree 1. In other words, the input are functions g_i for $i \in [k]$ where $g_i : [mb^i] \rightarrow [mb^{i+1}]$ for $i < k - 1$, and $g_{k-1} : [mb^{k-1}] \rightarrow [2]$. The next theorem implies that the communication is $\geq c_k m$.

Theorem 14.10. [218] Let P be a k -party one-way protocol using communication $\leq c_k mb$. Then $\mathbb{P}_x[P(x) \neq G_k^{m,b}] \geq c_k^m$.

For example, for $m = 1$ we require communication $\geq c_k b$, whereas $cb \log b$ suffices (it's the input length to the second party). Thus for fixed k this bound is nearly tight. (One can have a tight bound by considering trees, that is restricting the range of pointers; this slightly

complicates the exposition.) The total input length is $n \leq kb^{k-1}$. Thus for constant k one needs communication $\geq c_k n^{1/(k-1)}$. One can work out the dependence on k and show that the bound remains non-trivial for any $k \leq \log^c n$ where n is the total input length.

Proof. We proceed by induction on k . For every k we prove the statement for any setting of m, b . The base case $k = 1$ is clear: $P(x)$ is a fixed string, while G is a uniform string in $[2]^m$. The error probability is $\geq 2^{-m}$.

For the induction step, let P use communication t and $p := \mathbb{P}_x[P(x) \neq G_k^{m,b}]$. Write $x = (x_1, y)$ where x_1 is on the forehead of the first party. We have

$$\mathbb{P}_y \left[\mathbb{P}_{x_1}[P(x) \geq G_k^{m,b}] \geq p/2 \right] \geq p/2.$$

Let P_a be the protocol P where the first party always communicates string a , regardless of y . We claim there exists a s.t.

$$\mathbb{P}_y \left[\mathbb{P}_{x_1}[P_a(x) = G_k^{m,b}] \geq p/2 \right] \geq 2^{-t} p/2.$$

Note that the probability over x_1 is not reduced because the first party's message does not depend on x_1 .

Now let $m' := mb$ and define protocol P' for $G_{k-1}^{m',b}$. On input y , P' runs P_a for r times on inputs (x^i, y) for $i \in [r]$, where the x^i are independent choices for the first party's input. For any of the m' bits that are pointed to by some x^i , P' outputs the corresponding bit. In case different runs give different values, the answer can be arbitrary. For any bit that is not pointed to by any x^i , P' guesses at random. This gives a randomized protocol; one can fix the randomness and preserve the success probability.

The communication of P' is $\leq rt$.

To analyze the success probability. Fix any y for which $\mathbb{P}_{x_1}[P_a(x) = G_k^{m,b}] \geq p/2$. The probability that all the r runs are correct is $\geq (p/2)^r$. The probability that there are $\geq cm'$ bits that are not pointed to by some x^i is at most the probability that there is a set of size cm' s.t. mr pointers fall there, which is at most

$$\leq \binom{m'}{cm'} (1 - c)^{mr} \leq c^{m'} c^{-mr} \leq c^{m'},$$

for $r \geq cb$.

When that does not happen, the random guesses will be correct w.p. $\geq 2^{-cm'}$.

Overall, the success probability over uniform y is

$$\geq 2^{-t} p/2 \cdot ((p/2)^r - c^{m'}) \cdot 2^{-cm'}.$$

For $p \geq 2^{-cm}$ and $t \leq cm'$, the overall success probability is $\geq c^{m'}$. **QED**

In the case $k = 3$ according to Theorem 14.10 we need communication $\geq c\sqrt{n}$. As mentioned above, this is essentially tight for G , because of the way the layers are constructed. But one can consider the natural question of chasing pointers where each layer has n nodes. It is a tantalizing open question whether there is a protocol with communication $\leq c\sqrt{n}$. The trivial protocol takes communication n , and one might wonder if that's tight. But a clever protocol achieves sublinear communication.

14.3.3 Sublinear communication for 3 player

For $k = 3$ we consider pointer chasing on layers of sizes $1, n, n$. Define G as

$$G(i, g, h) := h(g(i))$$

where $i \in [n], g : [n] \rightarrow [n]$, and $h : [n] \rightarrow [2]$.

We consider the even more restricted *simultaneous* communication model where the players speak once, non-interactively. Naive intuition suggests that linear communication might be needed. In fact, such bounds were claimed several times, but each time the authors later realized that they only applied to special cases. Indeed, we have:

Theorem 14.11. [154, 39] PC has simultaneous communication $o(n)$.

Proof. We sketch the ideas in the proof [154] in case g is a permutation π . The case of general g is in [39]. Let H be a bipartite graph H between the n nodes in the middle layer and the n nodes in the last. For any permutation π , let $G_{H,\pi}$ denote the graph on the n last nodes where $\{x, y\}$ is an edge iff $\pi^{-1}(x)$ has an edge to y in H .

The main claim is that there is H of degree $d = (1 + \epsilon)pn$ s.t. for any π $G_{H,\pi}$ can be covered by $r = o(n)$ cliques. (Note any graph has a trivial covering with number of the edges cliques.)

The protocol is as follows. H is known to all.

Player 1, for each of the r cliques, announces the parity of the bits $h(x)$ for x in the clique.

Player 2 announces $h(x)$ for all the d neighbors x of i in H . This is d bits.

Player 3 knows $k := \pi(i)$. It considers the clique of $G_{H,\pi}$ containing k . It knows the parity of the $h(x)$ for x in this clique. Also, for any x in the clique, x and k are connected, hence $\pi^{-1}(k) = i$ is adjacent to x . So from the message of Player 2 we know $h(x)$. We can subtract off all these bits to get $h(k)$.

As stated, this protocol is not simultaneous. To make it simultaneous, let Player 3 announce which of the cliques k is in, and also which of the d neighbors of i are connected via H to nodes in that clique that are not k . Then the referee has a bit per clique, knows which bit to look at, and knows which bits of Player 2 to consider.

Player 3 message takes $\log r + d$.

The existence of H with suitable parameters can be established by the probabilistic method. Specifically, let H be distributed as $G(n, p)$, a random graph where each edge is present independently with prob. p . We observe that for any permutation $G_{H,\pi}$ is random from $G(n, p^2)$. Thus its complement is random from $G(n, 1 - p^2)$. One can show that w.h.p. this complement has chromatic number at most $r = o(n)$. It is known that the chromatic number of a graph is equal to the minimum number of independent sets that you need to cover the nodes of the graph. From this the result follows. **QED**

This protocol can be generalized to more players.

14.4 Notes

Several proofs of Theorem 14.3 exist [107, 160, 25], see the books [118, 156] or the survey [46] for two different expositions. For more on disjointness see also the survey [173].

Several presentations of the proof of Theorem 14.8 exist: [49, 157, 217]. We followed the latter.

A proof of Theorem 14.10 in the case $k = 3$ appeared in [20] but did not readily extend to larger k . The proof we presented is a streamlined version of the argument in [218]. The latter paper works with trees instead of graphs to obtain slightly better parameters at the cost of a slightly more involved analysis of the number of bits hit by pointers.

The simulation of Sym-And circuits by nof protocols () is from [86].

Chapter 15

Algebraic complexity

Stepping back, previous chapters have investigated the complexity of computing strings of length 2^n , corresponding to the truth-table of functions from $[2]^n$ to $[2]$ starting from basic strings (or functions) and changing them via simple operations. For example for boolean circuits the basic functions are the constant 0, 1, the variable x_i , and we combine them via And/Or/Not gates.

It is natural to consider other objects and to allow for different operations. In fact, we have already encountered other models which are more algebraic, like polynomials and matrices. In this chapter we explore more algebraic models, and in particular we consider computing other objects, namely polynomials.

Surprisingly, the development of this theory closely parallels that of its boolean counterpart. We will encounter again (especially starting in section §15.4) many of the main themes and results seen so far, including depth reduction, impossibility results for small-depth models that are “just short” of proving major separations, the grand challenge, reductions, completeness, and an algebraic analogue of NP.

15.1 Linear transformations, rigidity, and all that

tbd

15.2 Computing integers

Perhaps the simplest algebraic question is that of computing n -bit integers using arithmetic circuits over the integers, using no variables and no constants except 1 and -1 . Usual counting argument like that in the proof of Theorem 3.6 show that most n -bit integers require circuits of size $n/\log^c n$. And this is again nearly tight since any integer $t \in [2]^n$ can be computed with cn operations by writing $t = 2^0 t_0 + 2^1 t_1 + \dots + 2^{n-1} t_{n-1}$, and computing 2^i takes $\leq \log i + c$ operations via repeated squaring.

As usual, the grand challenge is to exhibit “explicit” integers that are hard to compute. In particular, integers that cannot be computed with $\log^c n$ operations. A prominent integer

in this context is the factorial. If it is easy, then factoring is also easy.

Theorem 15.1. Suppose n -bit factorial has algebraic circuits of size $\log^a n$. Then there are boolean circuits factoring n -bit integers of size n^{c_a} .

This connection is slightly more subtle than the similarity of the words “factorial” and “factoring” might indicate, and the difference in the circuit-size bounds might hint at its non-triviality. The idea is that the smallest t s.t. x divides $t!$ contains a non-trivial factor of x , which can be found computing the greatest common divisor.

Example 15.1. The following 12×12 matrix has $j! \bmod i$ in row i , column j . In a row, all entries to the right of a 0 are 0 as well. For $i > 4$, the position of the 0 gives a number with a non-trivial factor of i .

	1	2	3	4	5	6	7	8	9	10	11	12
1	0											
2	1	0										
3	1	2	0									
4	1	2	2	0								
5	1	2	1	4	0							
6	1	2	0									
7	1	2	6	3	1	6	0					
8	1	2	6	0								
9	1	2	6	6	3	0						
10	1	2	6	4	0							
11	1	2	6	2	10	5	2	5	1	10	0	
12	1	2	6	0								

Proof. We are given $x \in [2]^n$ that we’d like to factor. Suppose we can compute the smallest integer t s.t. x divides $t!$. Trivially, $t \leq x$. Also, if $x > 4$ is composite then $t < x$.

Exercise 15.1. Prove this. What if $x = 4$?

Let us then write $(t - 1)! = qx + r$ with $r < x$ and so $t! = qtx + rt$, and $x|rt$. Hence $sx = rt$ for an integer s . Because $t < x$ and $r < x$, the prime-power factors of x cannot be all in r or all in t . So the greatest common divisor of x and t is non-trivial, and we can compute it, divide x by it, and iterate. If the decomposition is trivial, that is $t = x$, then it means that x is prime by what we said above, and we can stop.

It remains how to determine t . We use binary search, beginning with $t = x$, noting that if $x|i!$ then $x|j!$ for any $j > i$. For each candidate t , we consider the algebraic circuit for $t!$.

We then use this circuit to compute $t! \bmod x$. Note that $t!$ has an unfeasible number of bits (exponential in n), but $t! \bmod x$ has $\leq n$ bits, so we can write it down – a classic move used extensively, e.g. in Exercise 2.20 and section 7.2.1.

By assumption, the algebraic circuit for $t!$ has size $\log^a(\log t!) \leq c_a \log^a t$. We have $t \leq x \leq 2^{n+1}$ and the result follows. **QED**

And we will see below in section §15.4 that if it is hard then another long-sought separation follows.

15.3 Univariate polynomials

A next natural question is computing univariate polynomials. An important distinction must be made. We can consider computing polynomials *formally*, which we can think of as a sequence of coefficients, or *informally*, as functions. This distinction disappears when the field is larger than the degree by Lemma 2.3, but otherwise leads to different theory. For example over \mathbb{F}_2 we have $x^2 = x$ informally (i.e. the identity holds for every field element) but obviously not formally. Obviously formal identities are also informal, so informal impossibility results are harder to establish than formal.

Given this, the cleanest setting may be when the underlying field is infinite.

—
Here we can allow arbitrary constants

The situation is similar to the previous section. A specific polynomial of interest is the approximation to the exponential function: $\sum_{i=0}^n X^i/i!$.

15.4 Multivariate polynomials

Again, the challenge is to exhibit “explicit” polynomials that are hard to compute.

For larger-depth there is a superlinear informal result that does not have a formal counterpart. For several explicit degree- d polynomials in n variables it can prove bounds of the form $cn \log d$. We state one example:

Theorem 15.2. [184, 27] Computing $\sum_{i \in [n]} x_i^d$ requires size $cn \log d$.

The proof is “simple and surprising:” it is shown in [27] that computing a polynomial p and simultaneously all its n partial derivatives w.r.t. the n variables only costs a constant factor more than computing p . Then one can use the “degree bound” from [184].

We now turn to constant-depth circuits.

Algebraic impossibility from boolean impossibility

Note that over the field \mathbb{F}_2 the informal imp. results obtained in section 10.3.1 (see especially Exercise 10.7) are algebraic (since And is like multiplication) – and nothing better is known even if one is informal, for small depth. The techniques in section 10.3.1 can be extended to slightly larger fields [76]. The latter paper gives exponential lower bounds for $\Sigma\Pi\Sigma$ circuits computing an explicit low-degree polynomial. The idea is similar to that in section 10.3.1: we show that such circuits are approximated by low-degree polynomials. We sketch why this is possible, to illustrate where the field size plays a role. It suffices to approximate $\Pi\Sigma$ circuits well. Consider one such circuit, and let r be the rank of the linear forms input to the Π gate (excluding their constants, if any). If the rank is large, then over a uniform input it’s likely that at least one linear form will be zero and so the whole circuit is zero. If the rank r is small, then we can write each linear form as a linear combination of $\leq r$ linear forms. Now if we expand the Π gate we will have sum of products of these r linear forms.

Now we can use the fact that over a field of size q we have $X^q = X$, so we reduce the degree of each form in any product to at most $q - 1$. Overall, the degree will be $\leq (q - 1)r$.

Large fields

But for larger, or infinite fields a different set of techniques appears necessary, and only formal results are known.

15.4.1 VNP

Similarly to NP, an important class of polynomials can be defined by summing over all boolean values of a set of variables.

Definition 15.1. The Σ -algebraic circuits $S(X_1, \dots, X_n)$ of size $\leq s$ are those that can be written as $\sum_{y_1, \dots, y_s \in \{0,1\}} C(X_1, \dots, X_n, y_1, \dots, y_s)$ where C is an algebraic circuit of size $\leq s$.

Several polynomials of interest that are not known to have small algebraic circuits can be shown to have small Σ -algebraic circuits.

Example 15.2. We show that the *permanent* polynomial in n variables,

$$p(x_0, x_1, \dots, x_{n-1}) = \sum_{\pi} \prod_{i \in [n]} x_{i, \pi(i)}$$

where π ranges over all permutations of $[n]$, has Σ -algebraic circuits of size n^c . It is an open problem whether it has (plain) arithmetic circuits of power size.

We will encode π using n^2 bits M specifying an $n \times n$ permutation matrix also written M . Suppose we have a polynomial g s.t. $g(M) = 1$ if M is a permutation and 0 otherwise. Then we have:

$$p = \sum_{M \in \{0,1\}^{n^2}} g(M) \prod_{i \in [n], j \in [n]} x_{i,j} \cdot M_{i,j}.$$

Thus it only remains to show that g has small algebraic circuits.

Exercise 15.2. Finish the example.

Similar to the P vs. NP question, the prominent question here is whether Σ -algebraic and algebraic circuits have similar power. The following unexpected result connects this question to the complexity of computing integers (section §15.2). One can get similar results for other “explicit” integers or even univariate polynomials (including those mentioned in 15.3).

Theorem 15.3. [115, 41] [If Σ -algebraic circuits are easy then so are univariate polynomials and integers] Suppose that every Σ -algebraic circuit in n variables of size s has an equivalent algebraic circuit of size s^d for some constant d . Then the 2^n -bit factorial has an algebraic circuit of size n^{cd} .

The proof is an excellent display of “scaling up and down” and connecting disparate complexity results.

Proof. First we claim bit i of the factorial given $i \in [2]^n$ is computable in

$$\text{Maj} \cdot \text{Maj} \cdot \dots \cdot \text{Maj} \cdot \text{PCkt},$$

where the number of applications of the Maj operator is c . (Cf. section §6.4 for the definition of the operator.) This follows from the fact that iterated multiplication of integers is in TC (Theorem 10.2).

Now note that for any function $f : [2]^n \rightarrow [2]$ in $\text{Maj} \cdot \text{PCkt}$ one can write down a Σ -algebraic circuit S of power size whose value on $x \in [2]^n$ is ≥ 0 iff $f(x) = 1$. (Note for $x, y \in [2]$ we have $x \wedge y = x \cdot y$ and $x \vee y = x + y - x \cdot y$ and $\neg x = 1 - x$; so the translation is immediate.)

By assumption, S has equivalent algebraic circuits of power size. Applying this c times to the above, we obtain an algebraic circuit C of power size such that $C(i)$ is bit i of the factorial.

Now consider

$$S'(X_{n-1}, \dots, X_0) := \sum_{j_0, j_1, \dots, j_{n-1} \in [2]} C(j) X_0^{j_0} X_1^{j_1} \cdot X_{n-1}^{j_{n-1}}$$

in the variables X_i . Note that $S'(2^{n-1}, \dots, 8, 4, 2, 1)$ equals the desired factorial. We can't immediately apply the hypothesis to S' , until we note

$$X_0^{j_0} = (X_0 j_0 + 1 - j_0)$$

which allows to write S' as a Σ -algebraic circuit. Applying the hypothesis again yields an equivalent power-size algebraic circuit C' , and then again the desired factorial is $C'(2^{n-1}, \dots, 8, 4, 2, 1)$ and the powers of 2 can be computed via repeated squaring. **QED**

15.5 Depth reduction in algebraic complexity

Various results are known. Work in the unbounded fan-in setting.

[197] shows that any algebraic circuit of size n^a computing a polynomial degree n^a has an equivalent circuit of size n^{c_a} and depth $c_a \log n$. This transformation does not have a counterpart in the boolean world.

Note however that this only works under the additional restriction that the circuit computes a low-degree polynomial:

Exercise 15.3. Give an algebraic circuit of size s that does not have an equivalent algebraic circuit of depth $\leq s$.

However, the following is similar in spirit to Theorem 9.6.

Theorem 15.4. Any n -variate polynomial of degree d computable by a size- n^a arithmetic circuit can be computed by a depth $\Sigma\Pi\Sigma$ of depth e and size $n^{c_a d^{1/(e-1)}}$.

In particular, for depth 3 the size is $n^{c_a \sqrt{d}}$. We shall see in section §15.7 that impossibility results for circuits of size $n^{c \sqrt{d}}$ computing degree- d polynomials are known. Thus, as mentioned at the beginning of the chapter, the situation in the algebraic world is strikingly analogous to that in the boolean world discussed in section §7.3. We have impossibility results for small-depth circuits that are “just short” of having major consequences for unbounded-depth models.

15.6 Completeness

The results in section §9.2, extended to other fields, show that iterated product of 3×3 matrices is complete for algebraic circuits. As in that section, the reduction is as simple as it gets: For any power-size circuit one can write down a power-size product where the matrix entries are either constants or variables that computes the same polynomial.

15.7 Impossibility results for small-depth circuits

Theorem 15.5. [123] Any depth- e circuit computing (one entry of) the product of d $m \times m$ matrices (a polynomial of degree d) has size $\geq n^{d^{c_e}}$, for any $d \leq \log n$. When $e = 3$ the bound is $n^{c \sqrt{d}}$.

This bound is tight.

This is a rapidly evolving area; subseq papers like Kush Saraf work with more intuitive partition of variables (i.e., random), and don’t involve set-multilinear.

The proof in [123] involves two main steps. In first step we show the following transformation:

Lemma 15.1. Any circuit of size s and depth d computing a set-multilinear polynomial has an equivalent set-multilinear circuit of depth $< 2d$ and size $s^c d^{cd}$, over a field of characteristic 0.

In the second step we prove impossibility for set-multilinear circuits.

The proof of Lemma 15.1 also involves two steps. The first step makes the fan-in of multiplication gates at most d . This makes the size blow up in the transformation to set multilinear tolerable. The first step is done by a generic step that makes the circuit homogeneous. For such circuits you can simply remove multiplication gates with fan-in larger than d .

15.8 Algebraic TMs

TBD

15.9 Notes

Impossibility results for univariate polynomials were first studied in [185]. That paper establishes negative results for polynomials with very large coefficients, but the results are non-trivial since arbitrary constants can be used by the circuit. For a survey see the book [42].


For an introduction to multivariate algebraic complexity see [42] and the survey [175].

The result in section §15.7 builds on a long line of works, see [124] for discussion.

Theorem 15.4 is the culmination of a line of works about depth reduction that was rekindled by [4]. See discussion of subsequent work in [77] for this statement for depth 3. The formulation for any depth is stated in [124].

Chapter 16

Barriers



In an attempt to understand the Grand Challenge (Chapter 3), one can identify several proof techniques and show or speculate that they cannot solve it. Such arguments are known as “barriers.” Several barriers have been put forth, and in fact there are even barriers to barriers, i.e., arguments indicating that proving a barrier is difficult, making complexity theory a rather philosophical and introspective field.

The two main barriers are the *black-box* (a.k.a. *oracle* or *relativization* barrier) and the *natural proofs* barrier.

16.1 Black-box

As hinted, many of the results we have shown don’t really exploit the specifics of the model we are working with, but work in greater generality. How to make this more precise? When programming, we can think of having access to a powerful *library*, or *subroutine*, or *oracle*, or *black-box*. Informally, if our argument extends in such cases we say it is black-box, or that it *relativizes*. Such a library is known as *oracle* or *black-box*. The black-box barrier helps us understand the limits of basic simulation arguments, including diagonalization, which tend to relativize.

Formally, we are going to equip our models, such as TMs, with access to a function $f : [2]^* \rightarrow [2]$, known as *oracle*. At any point, the model can ask for the value of the function at an input that has been computed. In the case of TMs, we can think of a special oracle tape and a special oracle state. Upon entering the state, the contents $x \in [2]^*$ on the oracle tape are replaced in one step with $f(x) \in [2]$. We can then define corresponding complexity classes, denoted P^f , and so on.

To illustrate, consider the separation $P \neq \text{Exp}$, which follows from the Time Hierarchy Theorem 3.3. The result relativizes:

Theorem 16.1. For every oracle $f : [2]^* \rightarrow [2]$, $P^f \neq \text{Exp}^f$.

Proof. Diagonalization and the time hierarchy work just as well for oracle machines. Specifically, when simulating a machine M running in time t with a machine M' running time $t' > t$, the simulation proceeds as before, and if M queries the oracle then M' does that as well. **QED**

Next we argue that relativizing techniques cannot resolve other major questions. Perhaps the simplest example is for P vs. PSpace , because the way the oracle is accessed is clear.

Claim 16.1. There is an oracle f s.t. $P^f = \text{PSpace}^f$.

Proof. The oracle takes as input $(M, 1^s, x)$, simulates M using space s on input x for $\leq |M|^s$ steps, and returns its answer. If M exceeds space s , the oracle returns 0. We claim that $\text{PSpace}^f = \text{PSpace}$. This is just because the oracle queries can be answered by direct simulation using power space. Further, $\text{PSpace} \subseteq P^f$. The proof is completed by combining the two claims. **QED**

Hence, relativizing proofs cannot prove the grand challenge that $P \neq PSpace$.

A less exciting possibility, ruled out next, is that relativizing techniques could show $P = PSpace$.

Claim 16.2. There is an oracle f s.t. $P^f \neq PSpace^f$.

The straightforward proof idea is to define f that always returns 0 on P machines, exploiting that they make few queries. Essentially, for such an oracle $P^f = P$. On the other hand, $PSpace$ machines, thanks to their ability of querying the oracle at every input of length n , will be able to compute non-trivial functions of f . Formalizing this idea requires defining the oracle in stages, to handle each P machine in turn, but this detail does not seem to further our understanding.

There is an industry of oracle constructions. As essentially seen in Chapter 6, PH^f basically corresponds to AC functions of the truth table of f , and one can use results about AC to give various oracle separations for PH and related classes.

Half a century ago oracle separations have enjoyed an exciting and prolific phase. A second phase has followed during which it was realized that oracles provide limited information and they were relegated as curiosities. More recently, they have made a comeback, sometimes under the new term of *black-box*, in cryptography and quantum computing, and seem to be experiencing the first phase again.

16.2 Natural proofs

The natural proofs barrier aims to explain the limit of *combinatorial* proof techniques. The idea is simple:

(1) Most combinatorial proof techniques against a class of functions F (for example, F are the functions on n bits computable by circuits of size n^{10} and depth 10) do more than providing a separation: They yield an *efficient* algorithm that given the truth table of length 2^n of a function can distinguish tables coming from F from those coming from uniformly random functions.

(2) The classes F for which we would like to prove impossibility are believed to be powerful enough to compute pseudorandom functions, i.e., truth tables that *cannot* be efficiently distinguished from uniformly random functions.

Note that (2) is not known unconditionally, but just believed to be the case. This is where complexity theory gets quite philosophical. We can't really *prove* (2) without solving the Grand Challenge, in which case these barriers are not actual barriers. On the other hand one can have a *belief* that (2) is indeed true even though we can't prove it, and if that's the case indeed to solve the Grand Challenge one needs to somehow bypass (1) and find alternative techniques. We currently don't seem to have such techniques.

16.2.1 TMs

We illustrate the natural-proofs barrier for 1TMs. Let us revisit the information bottleneck technique from section §3.1 to show that from it we can extract an efficient algorithm to distinguish truth-tables computed by fast 1TMs from uniform functions. One is tempted to consider a test checking if there is a large set where the function is constant, and moreover X is a product set $X = Y \times Z$. However, it is not clear that this test would be efficient. It is more convenient to use the simulation of 1TMs by low-communication protocols, Theorem 14.6, and use the quantity R from section 14.2.2.

16.2.1.1 Telling subquadratic-time 1TMs from random

Given the truth-table of a function $f : [2]^n \rightarrow [2]$, our test D will consider the function $f_0 : [2]^{n/3} \times [2]^{n/3} \rightarrow [2]$ defined as $f_0(x, y) := f(x0^{n/3}y)$, and check if $R(f_0) \geq 2^{-cn}$.

First, let us verify that fast 1TMs indeed pass D . Let M be an s -state 1TM running in time t computing $f : [2]^n \rightarrow [2]$. By Theorem 14.6, f_0 has 2-party protocols with communication $d := c(\log s)t/n$ and error $\leq 1/2$. By Lemma 14.2, $R(f_0) \geq 2^{d/c} \geq s^{ct/n}$.

Second, let's verify that D is efficiently computable. Indeed, following the definition we can compute D in time 2^{cn} , which is power in the input length 2^n .

Third, and finally, we show that random functions $U : [2]^n \rightarrow [2]$ don't pass the test. Indeed, we have

$$\mathbb{E}_U[R(U)] = \mathbb{E}_U \mathbb{E}_{\substack{x_1^0, x_2^0 \\ x_1^1, x_2^1}} [U(x_1^0, x_2^0) + U(x_1^0, x_2^1) + U(x_1^1, x_2^0) + U(x_1^1, x_2^1)].$$

When the x_1 are distinct and the x_2 are distinct, the expectation is 0. In the other case the expectation is 1, and this happens with prob. $\leq 2^{-cn}$. Therefore $\mathbb{E}_U[R(U)] \leq 2^{-cn}$. Moreover, R is never negative, for $R(f) = \mathbb{E}_{x_1^0} (\mathbb{E}_y e^{[f(x_1^0, y) + f(x_2^0, y)]})^2$. Hence

$$\mathbb{P}_U[R(U) \geq 2^{-cn}] \leq 2^{-cn}.$$

The upshot of the all the above is that we have devised an efficient test that can distinguish truth tables of functions computed by fast 1TMs from truth tables of uniformly random functions.

16.2.1.2 Quadratic-time 1TMs can compute pseudorandom functions

We now sketch a candidate pseudorandom function computable in quadratic time by 1TMs with cn^2 states. The candidate is an asymptotic generalization of a well-documented and widely used block cipher: the Advanced Encryption Standard, AES. AES is based on the substitution-permutation network (SPN) structure, and will actually compute a function from $[2]^n \rightarrow [2]^n$ (whereas range $[2]$ would suffice for our goals). On input $x \in [2]^n$, an SPN is computed over a number r of rounds, where each round “confuses” the input by dividing it into m/b bundles of b bits and applying a substitution function (S-box) to each bundle, and

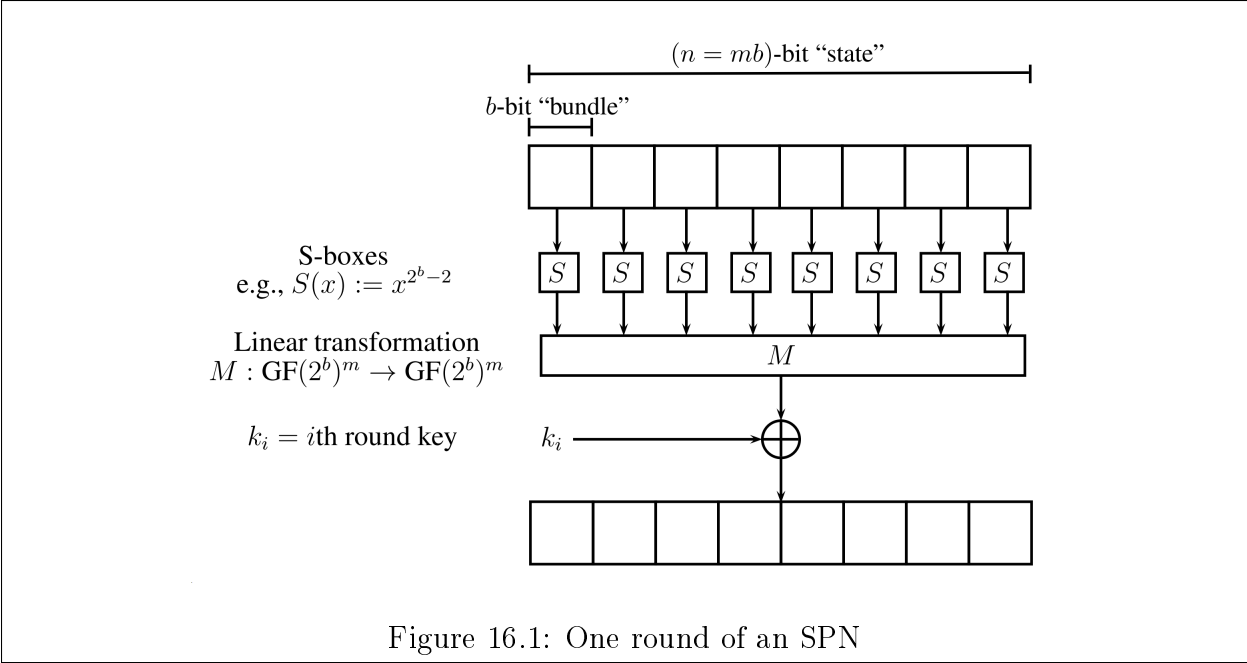


Figure 16.1: One round of an SPN

then “diffuses” the bundles by applying a matrix M with certain “branching” properties. At the end of each round i , the n bits are xor-ed with an n -bit round seed k_i , refer to figure 16.1.

The candidate follows the design considerations behind the AES block cipher, and particularly its S-box. For any n that is a multiple of 32, we break the input into $m := n/8$ bundles of $b = 8$ bits each, viewed as elements in the field \mathbb{F}_{2^8} , and perform $r = n$ rounds. We use the S-box $S(x) := x^{2^b-2}$. M is computed in two (linear) steps. In the first step, a permutation $\pi : [m] \rightarrow [m]$ is used to shuffle the b -bit bundles of the state; namely, bundle i moves to position $\pi(i)$. The permutation π is computed as follows. The m bundles are arranged into a $4 \times m/4$ matrix. Then row i of the matrix ($0 \leq i < 4$) is shifted circularly to the left by i places. In the second step, a maximal-branch-number matrix $\phi \in \mathbb{F}^{4 \times 4}$ is applied to each column of 4 bundles.

Let us now illustrate how one round can be computed in time cn with cn states. The bundles are written on the tape in column-major order: First the 4 bundles of the 1st column, then the 4 bundles of the 2nd column, and so on. The cn instances of S and φ can be computed in time cn . To see that π can also be computed in time cn , note that due to the representation, we can compute π with one pass, using that all but c bundles need to move $\leq c$ positions away. Finally, encoding the n -bit seed in the TM’s state transitions, the addition of each round key also takes time cn .

Therefore, the $r = n$ rounds can be computed in time cn^2 with cn^2 states.

By the simulation of TMs by circuits, this candidate is also computable by power-size circuits. A naive implementation gives fairly large depth, so next we consider smaller-depth circuits.

16.2.2 Small-depth circuits

In Chapter 10 we saw several impossibility results for AC. In the next exercise you are asked that at least one of the proof techniques we saw is natural.

Exercise 16.1. Give an efficient algorithm to distinguish truth tables of functions computable by power-size constant-depth AC from uniform.

Corollary 16.1. The ρ be a random restriction with n^{cd} stars. The probability that C_ρ is not a decision tree of depth $\log s$ is $\leq s2^{-n^{cd}}$.

There are candidate pseudorandom functions computable in TC. Some of them are linked to outstanding questions. For example, for one such candidate is known that if it is not pseudorandom than one can factor integers better than we currently know. The critical feature of TC that enables computing such candidates is iterated multiplication, see Theorem 10.2.

16.3 Notes

Relativization originated in the seminal work [24] and led to countless works on oracles. A variant of relativization where the oracles have additional algebraic structure is sharper for certain proof techniques and is studied in [1].

Natural proofs is from [159]. AES is described in [53]. The SPN structure of alternating “confusion” and “diffusion” steps was put forth already in [171]. The candidate 1TM pseudorandom function in section 16.2.1.2 is from [134].

The PRF in TC is from [139]. It gives TC of size $\geq n^c$. The work [10] considers TC of size $n^{1+\epsilon}$. [134] present a candidate, also based on AES, with these resources.

Chapter 17

I believe $P=NP$

“[...] Now it seems to me, however, to be completely within the realm of possibility that $\phi(n)$ grows that slowly. Since it seems that $\phi(n) \geq k \cdot n$ is the only estimation which one can obtain by a generalization of the proof of the undecidability of the Entscheidungsproblem and after all $\phi(n) \sim k \cdot n$ (or $\sim k \cdot n^2$) only means that the number of steps as opposed to trial and error can be reduced from N to $\log N$ (or $(\log N)^2$). However, such strong reductions appear in other finite problems [...]” [66]

(Cf. blog post [213])

The only things that matter in a theoretical study are those that you can prove, but it's always fun to speculate. After worrying about P vs. NP for half my life, and having carefully reviewed the available “evidence” I have decided I believe that $P = NP$.

A main justification for my belief is history:

1. In the 1950's Kolmogorov conjectured that multiplication of n -bit integers requires time $\geq cn^2$. That's the time it takes to multiply using the method that mankind has used for at least six millennia. Presumably, if a better method existed it would have been found already. Kolmogorov subsequently started a seminar where he presented again this conjecture. Within one week of the start of the seminar, Karatsuba discovered his famous algorithm running in time $cn^{\log_2 3} \approx n^{1.58}$. He told Kolmogorov about it, who became agitated and terminated the seminar. Karatsuba's algorithm unleashed a new age of fast algorithms, including the next one. I recommend Karatsuba's own account [108] of this compelling story.
2. In 1968 Strassen started working on proving that the standard cn^3 algorithm for multiplying two $n \times n$ matrices is optimal. Next year his landmark $cn^{\log_2 7} \approx n^{2.81}$ algorithm appeared in his paper “Gaussian elimination is not optimal” [183].
3. In the 1970s Valiant showed that the graphs of circuits computing certain linear transformations must be a *super-concentrator*, a graph which certain strong connectivity properties. He conjectured that super-concentrators must have a super-linear number of wires, from which super-linear circuit lower bounds follow [195]. However, he

later disproved the conjectured [196]: building on a result of Pinsker he constructed super-concentrators using a linear number of edges.

4. At the same time Valiant also defined *rigid* matrices and showed that an explicit construction of such matrices yields new circuit lower bounds. A specific matrix that was conjectured to be sufficiently rigid is the Hadamard matrix. Alman and Williams recently showed that, in fact, the Hadamard matrix is not rigid [11].
5. Constructing rigid matrices is one of *three* ways to get circuit lower bounds from a graph decomposition in [196]. Another way is via communication lower bounds. Here a specific candidate was the *sum-index* function, but then Sun [188] gave an efficient protocol for sum-index.
6. After finite automata, a natural step in lower bounds was to study slightly more general programs with constant memory. Consider a program that only maintains c bits of memory, and reads the input bits in a fixed order, where bits may be read several times. It seems quite obvious that such a program could not compute the majority function in polynomial time (see 1.1). This was explicitly conjectured by several people, including [37]. Barrington [135] famously disproved the conjecture by showing that in fact those seemingly very restricted constant-memory programs are in fact equivalent to log-depth circuits, which can compute majority (and many other things) (see Theorem 9.5).
7. Mansour, Nisan, and Tiwari conjectured [131] in 1990 that computing hash functions on n bits requires circuit size $\Omega(n \log n)$. Their conjecture was disproved in 2008 [103] where a circuit of size $O(n)$ was given.
8. For 30+ years the fastest run-time for graph isomorphism was exponential. A great deal was written on efficient proof systems for graph non-isomorphism. In 2015 Babai shocked the world with an almost power-time algorithm for graph isomorphism.
9. Maxflow is a central problem studied since the dawn of computer science. All solutions had running time $\geq n^{1+c}$, until a quasi-linear algorithm obtained in 2022.

And these are just some of the more famous ones. The list goes on and on. In number-on-forehead communication complexity, the function Majority-of-Majorities was a candidate for being hard for more than logarithmically many parties. This was disproved in [19] and subsequent works, where many other counter-intuitive protocols are presented, see section §14.3. Ditto for pointer chasing, see section 14.3.2. In data structures, would you think it possible to switch between binary and ternary representation of a number using constant time per digit and *zero* space overhead? Turns out it is [151, 55] (see section 12.1.2). Do you believe factoring is hard? Then you also believe there are pseudorandom generators where each output bit depends only on c input bits [15]. Known algorithms for directed connectivity use either super-polynomial time or polynomial memory. But if you are given access to polynomial memory full of junk that you can't delete, then you can solve directed connectivity using

only logarithmic (clean) memory and polynomial time [40]. And I haven't even touched on the many broken conjectures in cryptography, most recently related to obfuscation.

On the other hand, arguably the main thing that's surprising in the lower bounds we have is that they can be proved at all. The bounds themselves are hardly surprising. Of course, the issue may be that we can prove so few lower bounds that we shouldn't expect surprises. Some of the undecidability results I do consider surprising, for example Hilbert's 10th problem. But what is actually surprising in those results are the *algorithms*, showing that even very restricted models can simulate more complicated ones (same for the theory of NP completeness). In terms of lower bounds they all build on diagonalization, that is, go through every program and flip the answer, which is boring.

The evidence is clear: we have grossly underestimated the reach of efficient computation, in a variety of contexts. All signs indicate that we will continue to see bigger and bigger surprises in upper bounds, and $P=NP$. Do I really believe the formal inclusion $P=NP$? Maybe, let me not pick parameters. What I believe is that the idea that lower bounds are obviously true and we just can't prove them is not only baseless but even clashes with historical evidence. It's the upper bounds that are missing.

The “thousand different problems” argument for $P \neq NP$

“The class NP [...] contains thousands of different problems for which no efficient solving procedure is known.”[68]

“Among the NP-complete problems are many [...] for which serious effort has been expended on finding polynomial-time algorithms. Since either all or none of the NP-complete problems are in P, and so far none have been found to be in P, it is natural to conjecture that none are in P.” [95], Page 341.

I find these claims strange. In fact, the theory of NP completeness leads me to an opposite conclusion. As we saw, the problems can all be translated one into the other with extremely simple procedures, essentially doing *nothing*, just maybe complementing a bit. In what sense are they different? I think a good definition of different is that they are not known to be reducible to each other in a simple manner.

The “lots of people tried” argument for $P \neq NP$

The conjectures above were made by n top scientists in the area. On the other hand, $N \gg n$ people outside of the area attempted and failed to solve NP-hard problems. The fact that they are outsiders can be a strength or a weakness for the argument. It can be a strength, because of the sheer number, and because unshackled by the trends of the community, and without much interaction, the N people have been free to explore radically new ideas:

“Many of these problems have arisen in vastly different disciplines, and were the subject of extensive research by numerous different communities of scientists and engineers. These essentially independent studies have all failed to provide efficient algorithms for solving these problems, a failure that is extremely hard to attribute to sheer coincidence or a stroke of bad luck.”[68]

But it can also be a weakness, because unaware of the well-studied pitfalls, and with little communication, these N people are likely to all have followed the same route. Indeed, most of the countless bogus proofs claiming to resolve major open problem in complexity fail in one of only a handful of different ways. So it is likely that those N people don't quite count for N distinct attempts, but in fact a much smaller number, quite possibly less than n .

The “catastrophe” argument for $P \neq NP$

It's easy to consider scenarios in which $P = NP$ would not cause a catastrophe. A trivial scenario is if the algorithms take time n^d for exceedingly large d . A less obvious scenario is that the algorithms use complicated component X (think the classification of simple groups, or the 4-color theorem, etc.). And then we would enter a phase in which for a problem you ask if it can be solved without using X .

My “get stuck at the same point” argument for $P = NP$

The impossibility results that we have are sometimes obtained via seemingly very different proofs. For example, we have “bottom-up” and “top-down” proofs that AC can't compute parity. As far as I can tell, the proofs are genuinely different, I would argue more different than the various NP-complete problems (see the “thousand different problems” argument above). Why should different approaches stop at the same point, except because there is nothing else to prove? A related issue is why available techniques stop “right before” proving major results, see discussion e.g. in section §7.3. The most reasonable conclusion, it seems to me, is that this happens because the major results are actually false.

Throughout history, science has often proved wrong those who wouldn't take things at face value.

Complexity theory is perhaps unique in science. It appears that math is not ready for its problems. It is a bulwark against the business approach to science, the frenzy of the illusion of progress. For *ultimately* it doesn't matter how much you rake in or even who is writing bombastic recommendation letters for you, etc. These problems remain untouched. And progress may be more likely to come when you are alone, staring at blank paper:

“You do not need to leave your room. Remain sitting at your table and listen. Do not even listen, simply wait. Do not even wait, be quiet still and solitary. The world will freely offer itself to you to be unmasked, it has no choice, it will roll in ecstasy at your feet.” [106]

References

- [1] Scott Aaronson and Avi Wigderson. Algebrization: a new barrier in complexity theory. In *40th ACM Symp. on the Theory of Computing (STOC)*, pages 731–740, 2008.

- [2] Amir Abboud, Arturs Backurs, and Virginia Vassilevska Williams. Tight hardness results for LCS and other sequence similarity measures. In Venkatesan Guruswami, editor, *IEEE 56th Annual Symposium on Foundations of Computer Science, FOCS 2015, Berkeley, CA, USA, 17-20 October, 2015*, pages 59–78. IEEE Computer Society, 2015.
- [3] Leonard Adleman. Two theorems on random polynomial time. In *19th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 75–83. 1978.
- [4] Manindra Agrawal and V. Vinay. Arithmetic circuits: A chasm at depth four. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 67–75. IEEE Computer Society, 2008.
- [5] Miklós Ajtai. Σ_1^1 -formulae on finite structures. *Annals of Pure and Applied Logic*, 24(1):1–48, 1983.
- [6] Miklós Ajtai. A non-linear time lower bound for boolean branching programs. *Theory of Computing*, 1(1):149–176, 2005.
- [7] Miklos Ajtai and Avi Wigderson. Deterministic simulation of probabilistic constant-depth circuits. *Advances in Computing Research - Randomness and Computation*, 5:199–223, 1989.
- [8] Eric Allender. A note on the power of threshold circuits. In *30th Symposium on Foundations of Computer Science*, pages 580–584, Research Triangle Park, North Carolina, 30 October–1 November 1989. IEEE.
- [9] Eric Allender. The division breakthroughs. *Bulletin of the EATCS*, 74:61–77, 2001.
- [10] Eric Allender and Michal Koucký. Amplifying lower bounds by means of self-reducibility. *J. of the ACM*, 57(3), 2010.
- [11] Josh Alman and R. Ryan Williams. Probabilistic rank and matrix rigidity. In *ACM Symp. on the Theory of Computing (STOC)*, pages 641–652, 2017.
- [12] Noga Alon, László Babai, and Alon Itai. A fast and simple randomized algorithm for the maximal independent set problem. *Journal of Algorithms*, 7:567–583, 1986.
- [13] Noga Alon, Oded Goldreich, Johan Håstad, and René Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures & Algorithms*, 3(3):289–304, 1992.
- [14] Dana Angluin and Leslie G. Valiant. Fast probabilistic algorithms for hamiltonian circuits and matchings. *J. Comput. Syst. Sci.*, 18(2):155–193, 1979.
- [15] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0 . *SIAM J. on Computing*, 36(4):845–888, 2006.
- [16] Sanjeev Arora and Boaz Barak. *Computational Complexity*. Cambridge University Press, 2009. A modern approach.
- [17] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. of the ACM*, 45(3):501–555, May 1998.
- [18] László Babai. E-mail and the unexpected power of interaction. In *SCT*, pages 30–44. IEEE Computer Society, 1990.
- [19] László Babai, Anna Gál, Peter G. Kimmel, and Satyanarayana V. Lokam. Communi-

- cation complexity of simultaneous messages. *SIAM J. on Computing*, 33(1):137–166, 2003.
- [20] László Babai, Thomas P. Hayes, and Peter G. Kimmel. The cost of the missing bit: communication complexity with help. *Combinatorica. An Journal on Combinatorics and the Theory of Computing*, 21(4):455–488, 2001.
- [21] László Babai and Shlomo Moran. Arthur-merlin games: A randomized proof system, and a hierarchy of complexity classes. *J. Comput. Syst. Sci.*, 36(2):254–276, 1988.
- [22] László Babai, Noam Nisan, and Mária Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *J. of Computer and System Sciences*, 45(2):204–232, 1992.
- [23] Arturs Backurs and Piotr Indyk. Edit distance cannot be computed in strongly subquadratic time (unless SETH is false). *SIAM J. Comput.*, 47(3):1087–1097, 2018.
- [24] Theodore Baker, John Gill, and Robert Solovay. Relativizations of the $P=?NP$ question. *SIAM J. on Computing*, 4(4):431–442, 1975.
- [25] Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. of Computer and System Sciences*, 68(4):702–732, 2004.
- [26] Kenneth E. Batcher. Sorting networks and their applications. In *AFIPS Spring Joint Computing Conference*, volume 32, pages 307–314, 1968.
- [27] Walter Baur and Volker Strassen. The complexity of partial derivatives. *Theoret. Comput. Sci.*, 22(3):317–330, 1983.
- [28] Louay Bazzi. Polylogarithmic independence can fool DNF formulas. In *48th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 63–73, 2007.
- [29] Louay M. J. Bazzi. Polylogarithmic independence can fool DNF formulas. *SIAM J. Comput.*, 38(6):2220–2272, 2009.
- [30] Paul Beame. A switching lemma primer. Technical Report UW-CSE-95-07-01, Department of Computer Science and Engineering, University of Washington, November 1994. Available from <http://www.cs.washington.edu/homes/beame/>.
- [31] Paul Beame, Stephen A. Cook, and H. James Hoover. Log depth circuits for division and related problems. *SIAM J. Comput.*, 15(4):994–1003, 1986.
- [32] Paul Beame, Matei David, Toniann Pitassi, and Philipp Woelfel. Separating deterministic from nondeterministic multiparty communication complexity. In *34th Coll. on Automata, Languages and Programming (ICALP)*, pages 134–145. Springer, 2007.
- [33] Paul Beame, Michael Saks, Xiaodong Sun, and Erik Vee. Time-space trade-off lower bounds for randomized computation of decision problems. *J. of the ACM*, 50(2):154–195, 2003.
- [34] Richard Beigel and Jun Tarui. On ACC. *Computational Complexity*, 4(4):350–366, 1994.
- [35] Michael Ben-Or and Richard Cleve. Computing algebraic formulas using a constant number of registers. *SIAM J. on Computing*, 21(1):54–58, 1992.
- [36] Andrej Bogdanov and Emanuele Viola. Pseudorandom bits for polynomials. *SIAM J. on Computing*, 39(6):2464–2486, 2010.

- [37] Allan Borodin, Danny Dolev, Faith E. Fich, and Wolfgang J. Paul. Bounds for width two branching programs. In *Proceedings of the 15th Annual ACM Symposium on Theory of Computing, 25-27 April, 1983, Boston, Massachusetts, USA*, pages 87–93, 1983.
- [38] Mark Braverman. Poly-logarithmic independence fools AC^0 circuits. In *24th IEEE Conf. on Computational Complexity (CCC)*. IEEE, 2009.
- [39] Joshua Brody and Amit Chakrabarti. Sublinear communication protocols for multi-party pointer jumping and a related lower bound. In *25th Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 145–156, 2008.
- [40] Harry Buhrman, Richard Cleve, Michal Koucký, Bruno Loff, and Florian Speelman. Computing with a full memory: catalytic space. In *ACM Symp. on the Theory of Computing (STOC)*, pages 857–866, 2014.
- [41] Peter Bürgisser. On defining integers and proving arithmetic circuit lower bounds. *Comput. Complex.*, 18(1):81–103, 2009.
- [42] Peter Bürgisser, Michael Clausen, and Mohammad Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1997.
- [43] Samuel R. Buss and Ryan Williams. Limits on alternation trading proofs for time-space lower bounds. *Comput. Complex.*, 24(3):533–600, 2015.
- [44] J. Lawrence Carter and Mark N. Wegman. Universal classes of hash functions. *J. of Computer and System Sciences*, 18(2):143–154, 1979.
- [45] Ashok K. Chandra, Merrick L. Furst, and Richard J. Lipton. Multi-party protocols. In *15th ACM Symp. on the Theory of Computing (STOC)*, pages 94–99, 1983.
- [46] Arkadev Chattopadhyay and Toniann Pitassi. The story of set disjointness. *SIGACT News*, 41(3):59–85, 2010.
- [47] Lijie Chen and Roei Tell. Bootstrapping results for threshold circuits "just beyond" known lower bounds. In Moses Charikar and Edith Cohen, editors, *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, Phoenix, AZ, USA, June 23-26, 2019*, pages 34–41. ACM, 2019.
- [48] Benny Chor, Oded Goldreich, Johan Håstad, Joel Friedman, Steven Rudich, and Roman Smolensky. The bit extraction problem or t-resilient functions (preliminary version). In *26th Symposium on Foundations of Computer Science*, pages 396–407, Portland, Oregon, 21–23 October 1985. IEEE.
- [49] Fan R. K. Chung and Prasad Tetali. Communication complexity and quasi randomness. *SIAM Journal on Discrete Mathematics*, 6(1):110–123, 1993.
- [50] Richard Cleve. Towards optimal simulations of formulas by bounded-width programs. *Computational Complexity*, 1:91–105, 1991.
- [51] Stephen A. Cook. The complexity of theorem-proving procedures. In Michael A. Harrison, Ranan B. Banerji, and Jeffrey D. Ullman, editors, *Proceedings of the 3rd Annual ACM Symposium on Theory of Computing, May 3-5, 1971, Shaker Heights, Ohio, USA*, pages 151–158. ACM, 1971.
- [52] L. Csanky. Fast parallel matrix inversion algorithms. *SIAM J. Comput.*, 5(4):618–623,

- 1976.
- [53] Joan Daemen and Vincent Rijmen. *The Design of Rijndael: AES - The Advanced Encryption Standard*. Springer Verlag, 2002.
 - [54] Carsten Damm. Problems complete for $\oplus L$. *Inf. Process. Lett.*, 36(5):247–250, 1990.
 - [55] Yevgeniy Dodis, Mihai Pătraşcu, and Mikkel Thorup. Changing base without losing space. In *42nd ACM Symp. on the Theory of Computing (STOC)*, pages 593–602. ACM, 2010.
 - [56] Devdatt Dubhashi and Alessandro Panconesi. *Concentration of measure for the analysis of randomized algorithms*. Cambridge University Press, 2009.
 - [57] Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM J. Comput.*, 23(5):1001–1018, 1994.
 - [58] Lance Fortnow. Time-space tradeoffs for satisfiability. *J. Comput. Syst. Sci.*, 60(2):337–353, 2000.
 - [59] Aviezri S. Fraenkel and David Lichtenstein. Computing a perfect strategy for $n \times n$ chess requires time exponential in n . *J. Comb. Theory, Ser. A*, 31(2):199–214, 1981.
 - [60] Michael L. Fredman and Michael E. Saks. The cell probe complexity of dynamic data structures. In *ACM Symp. on the Theory of Computing (STOC)*, pages 345–354, 1989.
 - [61] Merrick L. Furst, James B. Saxe, and Michael Sipser. Parity, circuits, and the polynomial-time hierarchy. *Mathematical Systems Theory*, 17(1):13–27, 1984.
 - [62] Anka Gajentaan and Mark H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995.
 - [63] Anna Gál and Peter Bro Miltersen. The cell probe complexity of succinct data structures. *Theoretical Computer Science*, 379(3):405–417, 2007.
 - [64] M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
 - [65] Kurt Gödel. Über formal unentscheidbare sätze der principia mathematica und verwandter systeme, i. *Monatshefte für Mathematik und Physik*, 38:173–198, 1931.
 - [66] Kurt Gödel, 1956. Letter to John von Neumann. <https://www.anilada.com/notes/godel-letter.pdf>.
 - [67] Oded Goldreich. A sample of samplers - a computational perspective on sampling (survey). *Electronic Coll. on Computational Complexity (ECCC)*, 4(020), 1997.
 - [68] Oded Goldreich. *Computational Complexity: A Conceptual Perspective*. Cambridge University Press, 2008.
 - [69] Oded Goldreich. In a world of $p = bpp$. In *Studies in Complexity and Cryptography*, volume 6650 of *Lecture Notes in Computer Science*, pages 191–232. Springer, 2011.
 - [70] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity for all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):691–729, 1991.
 - [71] Oded Goldreich, Noam Nisan, and Avi Wigderson. On Yao’s XOR lemma. Technical Report TR95–050, *Electronic Colloquium on Computational Complexity*, March 1995. www.eccc.uni-trier.de/.

- [72] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In *40th ACM Symp. on the Theory of Computing (STOC)*, pages 113–122, 2008.
- [73] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.
- [74] Parikshit Gopalan, Raghu Meka, Omer Reingold, Luca Trevisan, and Salil Vadhan. Better pseudorandom generators from milder pseudorandom restrictions. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, 2012.
- [75] Raymond Greenlaw, H. James Hoover, and Walter Ruzzo. *Limits to Parallel Computation: P-Completeness Theory*. 02 2001.
- [76] Dima Grigoriev and Alexander A. Razborov. Exponential lower bounds for depth 3 arithmetic circuits in algebras of functions over finite fields. *Appl. Algebra Eng. Commun. Comput.*, 10(6):465–487, 2000.
- [77] Ankit Gupta, Pritish Kamath, Neeraj Kayal, and Ramprasad Satharishi. Arithmetic circuits: A chasm at depth 3. *SIAM J. Comput.*, 45(3):1064–1079, 2016.
- [78] Yuri Gurevich and Saharon Shelah. Nearly linear time. In *Logic at Botik, Symposium on Logical Foundations of Computer Science*, pages 108–118, 1989.
- [79] Dan Gutfreund and Emanuele Viola. Fooling parity tests with parity gates. In *8th Workshop on Randomization and Computation (RANDOM)*, pages 381–392. Springer, 2004.
- [80] Torben Hagerup. Fast parallel generation of random permutations. In *18th Coll. on Automata, Languages and Programming (ICALP)*, pages 405–416. Springer, 1991.
- [81] Kristoffer Arnsfelt Hansen, Oded Lachish, and Peter Bro Miltersen. Hilbert’s thirteenth problem and circuit complexity. In Yingfei Dong, Ding-Zhu Du, and Oscar H. Ibarra, editors, *Algorithms and Computation, 20th International Symposium, ISAAC 2009, Honolulu, Hawaii, USA, December 16-18, 2009. Proceedings*, volume 5878 of *Lecture Notes in Computer Science*, pages 153–162. Springer, 2009.
- [82] T. Hartman and R. Raz. On the distribution of the number of roots of polynomials and explicit weak designs. *Random Structures & Algorithms*, 23(3):235–263, 2003.
- [83] David Harvey and Joris van der Hoeven. Integer multiplication in time $O(n \log n)$. *Annals of Mathematics*, 193(2):563 – 617, 2021.
- [84] Johan Håstad. Almost optimal lower bounds for small depth circuits. In Juris Hartmanis, editor, *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 6–20. ACM, 1986.
- [85] Johan Håstad. *Computational limitations of small-depth circuits*. MIT Press, 1987.
- [86] Johan Håstad and Mikael Goldmann. On the power of small-depth threshold circuits. *Computational Complexity*, 1(2):113–129, 1991.
- [87] Pooya Hatami and William Hoza. Theory of unconditional pseudorandom generators. *Electron. Colloquium Comput. Complex.*, TR23-019, 2023.
- [88] Alexander Healy, Salil P. Vadhan, and Emanuele Viola. Using nondeterminism to amplify hardness. *SIAM J. on Computing*, 35(4):903–931, 2006.
- [89] F. C. Hennie. Crossing sequences and off-line turing machine computations. In *Symposium on Switching Circuit Theory and Logical Design (SWCT) (FOCS)*, pages 168–172,

- 1965.
- [90] F. C. Hennie. One-tape, off-line turing machine computations. *Information and Control*, 8(6):553–578, 1965.
 - [91] Fred Hennie and Richard Stearns. On the computational complexity of algorithms. *Transactions of the American Mathematical Society*, 117:285–306, 1965.
 - [92] Fred Hennie and Richard Stearns. Two-tape simulation of multitape turing machines. *J. of the ACM*, 13:533–546, October 1966.
 - [93] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4):439–561 (electronic), 2006.
 - [94] John E. Hopcroft, Wolfgang J. Paul, and Leslie G. Valiant. On time versus space. *J. ACM*, 24(2):332–337, 1977.
 - [95] John E. Hopcroft and Jeffrey D. Ullman. *Formal languages and their relation to automata*. Addison-Wesley Longman Publishing Co., Inc., 1969.
 - [96] John T. Gill III. Computational complexity of probabilistic turing machines. In *STOC*, pages 91–95. ACM, 1974.
 - [97] Neil Immerman. Nondeterministic space is closed under complementation. *SIAM J. Comput.*, 17(5):935–938, 1988.
 - [98] Russell Impagliazzo. Hard-core distributions for somewhat hard problems. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 538–545, 1995.
 - [99] Russell Impagliazzo and Ramamohan Paturi. The complexity of k -sat. In *IEEE Conf. on Computational Complexity (CCC)*, pages 237–, 1999.
 - [100] Russell Impagliazzo, Ramamohan Paturi, and Michael E. Saks. Size-depth tradeoffs for threshold circuits. *SIAM J. Comput.*, 26(3):693–707, 1997.
 - [101] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Computer & Systems Sciences*, 63(4):512–530, Dec 2001.
 - [102] Russell Impagliazzo and Avi Wigderson. $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *29th ACM Symp. on the Theory of Computing (STOC)*, pages 220–229. ACM, 1997.
 - [103] Yuval Ishai, Eyal Kushilevitz, Rafail Ostrovsky, and Amit Sahai. Cryptography with constant computational overhead. In *40th ACM Symp. on the Theory of Computing (STOC)*, pages 433–442, 2008.
 - [104] Hamid Jahanjou, Eric Miles, and Emanuele Viola. Local reductions. *Information and Computation*, 261(2), 2018. Available at <http://www.ccs.neu.edu/home/viola/>.
 - [105] Stasys Jukna. *Boolean Function Complexity: Advances and Frontiers*. Springer, 2012.
 - [106] Franz Kafka. Aphorisms.
 - [107] Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992.
 - [108] A. A. Karatsuba. The complexity of computations. *Trudy Mat. Inst. Steklov.*, 211(Optim. Upr. i Differ. Uravn.):186–202, 1995.
 - [109] Richard M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, The IBM Research Symposia Series, pages 85–103. Plenum

- Press, New York, 1972.
- [110] Richard M. Karp and Richard J. Lipton. Turing machines that take advice. *L'Enseignement Mathématique. Revue Internationale. IIe Série*, 28(3-4):191–209, 1982.
 - [111] Zander Kelley, Shachar Lovett, and Raghu Meka. Explicit separations between randomized and deterministic number-on-forehead communication. *CoRR*, abs/2308.12451, 2023.
 - [112] Adam Klivans and Rocco A. Servedio. Boosting and hard-core sets. *Machine Learning*, 53(3):217–238, 2003.
 - [113] Adam R. Klivans and Dieter van Melkebeek. Graph nonisomorphism has subexponential size proofs unless the polynomial-time hierarchy collapses. In *ACM Symposium on Theory of Computing (Atlanta, GA, 1999)*, pages 659–667. ACM, New York, 1999.
 - [114] Kojiro Kobayashi. On the structure of one-tape nondeterministic turing machine time hierarchy. *Theor. Comput. Sci.*, 40:175–193, 1985.
 - [115] Pascal Koiran. Valiant’s model and the cost of computing integers. *Comput. Complex.*, 13(3-4):131–146, 2005.
 - [116] Kenneth Krohn, W. D. Maurer, and John Rhodes. Realizing complex Boolean functions with simple groups. *Information and Control*, 9:190–195, 1966.
 - [117] S.-Y. Kuroda. Classes of languages and linear-bounded automata. *Inf. Control.*, 7(2):207–223, 1964.
 - [118] Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.
 - [119] Kasper Green Larsen. Higher cell probe lower bounds for evaluating polynomials. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 293–301, 2012.
 - [120] Kasper Green Larsen, Omri Weinstein, and Huacheng Yu. Crossing the logarithmic barrier for dynamic boolean data structure lower bounds. *SIAM J. Comput.*, 49(5), 2020.
 - [121] Leonid A. Levin. Universal sequential search problems. *Problemy Peredachi Informat-sii*, 9(3):115–116, 1973.
 - [122] Rudolf Lidl and Harald Niederreiter. *Finite fields*, volume 20 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, second edition, 1997.
 - [123] Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Superpolynomial lower bounds against low-depth algebraic circuits. In *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*, pages 804–814. IEEE, 2021.
 - [124] Nutan Limaye, Srikanth Srinivasan, and Sébastien Tavenas. Guest column: Lower bounds against constant-depth algebraic circuits. *SIGACT News*, 53(2):40–62, 2022.
 - [125] Nathan Linial and Noam Nisan. Approximate inclusion-exclusion. *Combinatorica*, 10(4):349–365, 1990.
 - [126] Shachar Lovett. Unconditional pseudorandom generators for low degree polynomials. In *40th ACM Symp. on the Theory of Computing (STOC)*, pages 557–562, 2008.
 - [127] Chi-Jen Lu, Shi-Chun Tsai, and Hsin-Lung Wu. Improved hardness amplification in NP. *Theor. Comput. Sci.*, 370(1-3):293–298, 2007.

- [128] Carsten Lund, Lance Fortnow, Howard Karloff, and Noam Nisan. Algebraic methods for interactive proof systems. *J. of the ACM*, 39(4):859–868, October 1992.
- [129] O. B. Lupanov. A method of circuit synthesis. *Izv. VUZ Radiofiz.*, 1:120–140, 1958.
- [130] Wolfgang Maass and Amir Schorr. Speed-up of Turing machines with one work tape and a two-way input tape. *SIAM J. on Computing*, 16(1):195–202, 1987.
- [131] Yishay Mansour, Noam Nisan, and Prason Tiwari. The computational complexity of universal hashing. *Theoretical Computer Science*, 107:121–133, 1993.
- [132] Yossi Matias and Uzi Vishkin. Converting high probability into nearly-constant time-with applications to parallel hashing. In *23rd ACM Symp. on the Theory of Computing (STOC)*, pages 307–316, 1991.
- [133] Pierre McKenzie and Stephen A. Cook. The parallel complexity of abelian permutation group problems. *SIAM J. Comput.*, 16(5):880–909, 1987.
- [134] Eric Miles and Emanuele Viola. Substitution-permutation networks, pseudorandom functions, and natural proofs. *J. of the ACM*, 62(6), 2015.
- [135] David A. Mix Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC^1 . *J. of Computer and System Sciences*, 38(1):150–164, 1989.
- [136] Cristopher Moore and Stephan Mertens. *The Nature of Computation*. Oxford University Press, 2011.
- [137] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. In *22nd ACM Symp. on the Theory of Computing (STOC)*, pages 213–223. ACM, 1990.
- [138] Joseph Naor and Moni Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM J. on Computing*, 22(4):838–856, 1993.
- [139] Moni Naor and Omer Reingold. Number-theoretic constructions of efficient pseudorandom functions. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 458–467, 1997.
- [140] E. I. Nechiporuk. A boolean function. *Soviet Mathematics-Doklady*, 169(4):765–766, 1966.
- [141] Valery A. Nepomnjaščii. Rudimentary predicates and Turing calculations. *Soviet Mathematics-Doklady*, 11(6):1462–1465, 1970.
- [142] NEU. From RAM to SAT. Available at <http://www.ccs.neu.edu/home/viola/>, 2012.
- [143] Noam Nisan. Pseudorandom bits for constant depth circuits. *Combinatorica. An Journal on Combinatorics and the Theory of Computing*, 11(1):63–70, 1991.
- [144] Noam Nisan. The communication complexity of threshold gates. In *Combinatorics, Paul Erdős is Eighty, number 1 in Bolyai Society Mathematical Studies*, pages 301–315, 1993.
- [145] Noam Nisan and Avi Wigderson. Hardness vs randomness. *J. of Computer and System Sciences*, 49(2):149–167, 1994.
- [146] Ryan O’Donnell. Hardness amplification within NP . *J. of Computer and System Sciences*, 69(1):68–94, August 2004.
- [147] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.

- [148] Christos H. Papadimitriou. *Computational Complexity*. Addison–Wesley, 1994.
- [149] Christos H. Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. Syst. Sci.*, 43(3):425–440, 1991.
- [150] Christos H. Papadimitriou and Stathis Zachos. Two remarks on the power of counting. In *Theoretical Computer Science*, volume 145 of *Lecture Notes in Computer Science*, pages 269–276. Springer, 1983.
- [151] Mihai Pătrașcu. Succincter. In *49th IEEE Symp. on Foundations of Computer Science (FOCS)*. IEEE, 2008.
- [152] Wolfgang J. Paul, Nicholas Pippenger, Endre Szemerédi, and William T. Trotter. On determinism versus non-determinism and related problems (preliminary version). In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 429–438, 1983.
- [153] Nicholas Pippenger and Michael J. Fischer. Relations among complexity measures. *J. of the ACM*, 26(2):361–381, 1979.
- [154] Pavel Pudlák, Vojtěch Rödl, and Jiří Sgall. Boolean circuits, tensor ranks, and communication complexity. *SIAM J. on Computing*, 26(3):605–633, 1997.
- [155] C. Radhakrishna Rao. Factorial experiments derivable from combinatorial arrangements of arrays. *Suppl. J. Roy. Statist. Soc.*, 9:128–139, 1947.
- [156] Anup Rao and Amir Yehudayoff. *Communication complexity*. 2019. <https://homes.cs.washington.edu/~anuprao/pubs/book.pdf>.
- [157] Ran Raz. The BNS-Chung criterion for multi-party communication complexity. *Computational Complexity*, 9(2):113–122, 2000.
- [158] Alexander Razborov. Lower bounds on the dimension of schemes of bounded depth in a complete basis containing the logical addition function. *Akademiya Nauk SSSR. Matematicheskie Zametki*, 41(4):598–607, 1987. English translation in *Mathematical Notes of the Academy of Sci. of the USSR*, 41(4):333–338, 1987.
- [159] Alexander Razborov and Steven Rudich. Natural proofs. *J. of Computer and System Sciences*, 55(1):24–35, August 1997.
- [160] Alexander A. Razborov. On the distributional complexity of disjointness. *Theor. Comput. Sci.*, 106(2):385–390, 1992.
- [161] Alexander A. Razborov. A simple proof of Bazzi’s theorem. *ACM Transactions on Computation Theory (TOCT)*, 1(1), 2009.
- [162] Omer Reingold. Undirected connectivity in log-space. *J. of the ACM*, 55(4), 2008.
- [163] J. M. Robson. N by N checkers is exptime complete. *SIAM J. Comput.*, 13(2):252–267, 1984.
- [164] J. M. Robson. An $O(T \log T)$ reduction from RAM computations to satisfiability. *Theoretical Computer Science*, 82(1):141–149, 1991.
- [165] Eyal Rozenman and Salil P. Vadhan. Derandomized squaring of graphs. In *APPROX-RANDOM*, pages 436–447, 2005.
- [166] Rahul Santhanam. On separators, segregators and time versus space. In *IEEE Conf. on Computational Complexity (CCC)*, pages 286–294, 2001.
- [167] Walter J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4(2):177–192, 1970.

- [168] M. Schaefer and C. Umans. Completeness in the polynomial-time hierarchy: a compendium. *SIGACT News, Complexity Theory Column*, 2002.
- [169] Arnold Schönhage. Storage modification machines. *SIAM J. Comput.*, 9(3):490–508, 1980.
- [170] Adi Shamir. $IP = PSPACE$. *J. of the ACM*, 39(4):869–877, October 1992.
- [171] Claude Shannon. Communication theory of secrecy systems. *Bell Systems Technical Journal*, 28(4):656–715, 1949.
- [172] Claude E. Shannon. The synthesis of two-terminal switching circuits. *Bell System Tech. J.*, 28:59–98, 1949.
- [173] Alexander A. Sherstov. Communication complexity theory: Thirty-five years of set disjointness. In *Symp. on Math. Foundations of Computer Science (MFCS)*, pages 24–43, 2014.
- [174] Victor Shoup. New algorithms for finding irreducible polynomials over finite fields. *Mathematics of Computation*, 54(189):435–447, 1990.
- [175] Amir Shpilka and Amir Yehudayoff. Arithmetic circuits: A survey of recent results and open questions. *Found. Trends Theor. Comput. Sci.*, 5(3-4):207–388, 2010.
- [176] Alan Siegel. On universal classes of extremely random constant-time hash functions. *SIAM J. on Computing*, 33(3):505–543, 2004.
- [177] Michael Sipser. A complexity theoretic approach to randomness. In *ACM Symp. on the Theory of Computing (STOC)*, pages 330–335, 1983.
- [178] Michael Sipser. *Introduction to the theory of computation, 3rd ed.* PWS Publishing Company, 1997.
- [179] Roman Smolensky. Algebraic methods in the theory of lower bounds for Boolean circuit complexity. In *19th ACM Symp. on the Theory of Computing (STOC)*, pages 77–82. ACM, 1987.
- [180] Roman Smolensky. On representations by low-degree polynomials. In *34th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 130–138, 1993.
- [181] Larry Stockmeyer and Albert R. Meyer. Cosmological lower bound on the circuit complexity of a small problem in logic. *J. ACM*, 49(6):753–784, 2002.
- [182] Larry J. Stockmeyer. The polynomial-time hierarchy. *Theor. Comput. Sci.*, 3(1):1–22, 1976.
- [183] Volker Strassen. Gaussian elimination is not optimal. *Numer. Math.*, 13:354–356, 1969.
- [184] Volker Strassen. Die rechnerischkomplexität von elementarsymmetrischen funktionen und von interpolationskoefizienten. *Numer. Math.*, 20:238–251, 1973.
- [185] Volker Strassen. Polynomials with rational coefficients which are hard to compute. *SIAM J. Comput.*, 3:128–149, 1974.
- [186] B. A. Subbotovskaya. Realizations of linear functions by formulas using $+$, $*$, $-$. *Soviet Mathematics-Doklady*, 2:110–112, 1961.
- [187] Madhu Sudan, Luca Trevisan, and Salil Vadhan. Pseudorandom generators without the XOR lemma. *J. of Computer and System Sciences*, 62(2):236–266, 2001.
- [188] Xiaoming Sun. A 3-party simultaneous protocol for SUM-INDEX. *Algorithmica*, 36(1):89–111, 2003.

- [189] Róbert Szelepcsényi. The method of forcing for nondeterministic automata. *Bull. EATCS*, 33:96–99, 1987.
- [190] Justin Thaler. Proofs, arguments, and zero-knowledge. *Found. Trends Priv. Secur.*, 4(2-4):117–660, 2022.
- [191] Seinosuke Toda. PP is as hard as the polynomial-time hierarchy. *SIAM J. on Computing*, 20(5):865–877, 1991.
- [192] Luca Trevisan. The program-enumeration bottleneck in average-case complexity theory. In *CCC*, pages 88–95. IEEE Computer Society, 2010.
- [193] Alan M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proc. London Math. Soc.*, s2-42(1):230–265, 1937.
- [194] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.
- [195] Valiant. On non-linear lower bounds in computational complexity. In *ACM Symp. on the Theory of Computing (STOC)*, pages 45–53, 1975.
- [196] Leslie G. Valiant. Graph-theoretic arguments in low-level complexity. In *6th Symposium on Mathematical Foundations of Computer Science*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977.
- [197] Leslie G. Valiant, Sven Skyum, S. Berkowitz, and Charles Rackoff. Fast parallel computation of polynomials using few processors. *SIAM J. Comput.*, 12(4):641–644, 1983.
- [198] Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.
- [199] J. H. van Lint. *Introduction to coding theory*. Springer-Verlag, Berlin, third edition, 1999.
- [200] Dieter van Melkebeek. A survey of lower bounds for satisfiability and related problems. *Foundations and Trends in Theoretical Computer Science*, 2(3):197–303, 2006.
- [201] Dieter van Melkebeek and Ran Raz. A time lower bound for satisfiability. *Theor. Comput. Sci.*, 348(2-3):311–320, 2005.
- [202] N. V. Vinodchandran. A note on the circuit complexity of PP. *Theor. Comput. Sci.*, 347(1-2):415–418, 2005.
- [203] Emanuele Viola. New lower bounds for probabilistic degree and AC0 with parity gates. *Theory of Computing*. Available at <http://www.ccs.neu.edu/home/viola/>.
- [204] Emanuele Viola. The complexity of constructing pseudorandom generators from hard functions. *Computational Complexity*, 13(3-4):147–188, 2004.
- [205] Emanuele Viola. Gems of theoretical computer science. Lecture notes of the class taught at Northeastern University. Available at <http://www.ccs.neu.edu/home/viola/classes/gems-08/index.html>, 2009.
- [206] Emanuele Viola. On approximate majority and probabilistic time. *Computational Complexity*, 18(3):337–375, 2009.
- [207] Emanuele Viola. On the power of small-depth computation. *Foundations and Trends in Theoretical Computer Science*, 5(1):1–72, 2009.
- [208] Emanuele Viola. The sum of d small-bias generators fools polynomials of degree d . *Computational Complexity*, 18(2):209–217, 2009.

- [209] Emanuele Viola. Reducing 3XOR to listing triangles, an exposition. Available at <http://www.ccs.neu.edu/home/viola/>, 2011.
- [210] Emanuele Viola. Bit-probe lower bounds for succinct data structures. *SIAM J. on Computing*, 41(6):1593–1604, 2012.
- [211] Emanuele Viola. The complexity of distributions. *SIAM J. on Computing*, 41(1):191–218, 2012.
- [212] Emanuele Viola. The communication complexity of addition. *Combinatorica*, pages 1–45, 2014.
- [213] Emanuele Viola, 2018. <https://emanueleviola.wordpress.com/2018/02/16/i-believe-pnp/>.
- [214] Emanuele Viola. Lower bounds for data structures with space close to maximum imply circuit lower bounds. *Theory of Computing*, 15:1–9, 2019. Available at <http://www.ccs.neu.edu/home/viola/>.
- [215] Emanuele Viola. Non-abelian combinatorics and communication complexity. *SIGACT News, Complexity Theory Column*, 50(3), 2019.
- [216] Emanuele Viola. Pseudorandom bits and lower bounds for randomized turing machines. *Theory of Computing*, 18(10):1–12, 2022.
- [217] Emanuele Viola and Avi Wigderson. Norms, XOR lemmas, and lower bounds for polynomials and protocols. *Theory of Computing*, 4:137–168, 2008.
- [218] Emanuele Viola and Avi Wigderson. One-way multiparty communication lower bound for pointer jumping with applications. *Combinatorica*, 29(6):719–743, 2009.
- [219] Avi Wigderson. *Mathematics and Computation: A Theory Revolutionizing Technology and Science*. Princeton University Press, 2019.
- [220] Ryan Williams. *Algorithms and Resource Requirements for Fundamental Problems*. PhD thesis, Carnegie Mellon University, 2007.
- [221] Ryan Williams. Nonuniform ACC circuit lower bounds. *J. of the ACM*, 61(1):2:1–2:32, 2014.
- [222] Andrew Yao. Theory and applications of trapdoor functions. In *23rd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 80–91. IEEE, 1982.
- [223] Andrew Yao. Separating the polynomial-time hierarchy by oracles. In *26th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 1–10, 1985.
- [224] Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity (extended abstract). In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 222–227. IEEE Computer Society, 1977.
- [225] Andrew Chi-Chih Yao. On ACC and threshold circuits. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 619–627, 1990.

“All that he does seems to him, it is true, extraordinarily new, but also, because of the incredible spate of new things, extraordinarily amateurish, indeed scarcely tolerable, incapable of becoming history, breaking short the chain of the generations, cutting off for the first time at its most profound source the music of the world, which before him could at least be divined. Sometimes in his arrogance he has more anxiety for the world than for himself.” [106]

Chapter 18

Annotated meta-bibliography

TBD

[105] Impossibility results in complexity theory.

[122] Massive reference for finite fields, though the focus is not computational.

[147]

[156]

[56]

[219]

[42]

[194]

[87]

[16] An excellent reference book covering for the first time more recent material, including the PCP theorem.

[68] A book providing a conceptual and personal perspective on complexity, also covering for the first time more recent material. Somewhat narrower in terms of topics compared to [16].

[95] An enticing masterpiece, still a great reference. Stay away from the watered-down “revisions.”

[148] Still a masterwork after 30+ years, conveying passion and originality, dated as it is in terms of material. Seferis’ poem is worth quoting in full:

*I wish nothing else but to speak simply
please grant me this privilege
because we have burdened our song with so much music
that it is slowly sinking
and our art has become so ornate
that the makeup has corroded her face
and it is time to say our few simple words
because tomorrow our soul sails away*

[Logicomix]: Another very enjoyable read. Wittgenstein was hilarious, but then again he always is. But if you like me were expecting an Armageddon or a Ragnarok unleashed by

Godel's incompleteness, you are going to be disappointed. There are only *c* scenes with him, and while one does show him on the board scribbling actual lines from his masterwork, the meaning and firepower of his discovery is nowhere to be found in the book, which ends instead with yet another...

[Stoner], by John Williams: Highly recommended, especially if you wonder about the meaning of life, especially academic life, or even marital life.

[The first world war, the second world war], by John Keegan. Masterful accounts of key events that shaped the geography of science production.

Chapter 19

Cryptography

crypto in NC0 (with Barrington?)

19.1 Natural proofs

obfuscation TMs

Index

- 3CNF, 51
- 3Color, 55
- 3Cycle-Detection, 50
- 3Sat, 51
- 3Sum, 49
- 4-Color, 59

- A**
- AC, 22
- AC[m], 131
- ACC, 131
- alphabet size*, 19
- AltCkt, 22
- alternating circuit, 22
- Alternation, 81
- AM-GM inequality, 28
- arithmetic circuit, 31
- arithmetic-mean geometric mean inequality, 28

- B**
- biased RRAM, 29
- BIG, 158
- bounded-intersection generator*, 158
- BPP, 27
- BPTIME, 27
- branching program, 98
- breaks, 151
- brute-force, 29

- C**
- Chernoff bound, Theorem 2.8, 27
- circuit, 22
- circuit hierarchy, 44
- Ckt, 22
- clause (in a CNF), 22

- Clique, 51
- CNF, 22
- collapse of PH, 82
- Collinearity, 49
- Compare-Exchange, 72
- computability thesis, 16
- computation table, 23
- configuration, 15
- configuration graph*, 99
- Cook-Levin theorem, Theorem 5.2, 68
- correlation*, 127
- cosmological results, 44
- crossing sequence*, 37
- CS, 37

- D**
- D , 27
- depth-reduction*, 65
- derandomization, 33
- deviation bounds, 27
- diagonalization*, 40
- direct addressing, 24
- distinguishes, 151
- divergence, 27
- DNF, 22

- E**
- error reduction, 28
- error-correcting code, 33
- ETH, 51
- Exp, 18
- expander graphs*, 153
- Exp-completeness, 77
- expected running time, 29
- Exponential time hypothesis, 51

F

\mathbb{F} , 30
 fan-in, 22
 fan-out, 22
 finite field, 30
 fool, 151
fools, 151
 Fourier expansion, Exercise 13.6, 154

G

Gap-3Sat, 62
 Gap-Maj, 85
 Generality, 14

H

hardcore-set, 161
hashing, 78
 HIT, 162
 hitter, 162
 hybrid method, 156

I

impossibility results, 36
inapproximability, 62
 information bottleneck, 37
input length, 18

L

L, 96
 linear time, 43
 local, 48
 Locality, 14
 logic, 44

M

MAlloc, 25
map reduction, 47
 Markov's inequality, Claim 2.1, 27
 Max-3Sat, 62, 77
 mergesort, 72
 Min-Ckt, 81
 MTM, 20
 Multiplication, 48
 Multi-tape machines, 20

N

NExp, 66
 NExp completeness, 76
 non-boolean outputs, 18
 Nondeterministic computation, 66
 NP, 66
 NP-complete, 68
 NP-completeness, 68
 NP-hard, 68
 NTime, 66

O

oblivious, 72
 oblivious TM, 20
 Odd-Even-Mergesort, 72
 Or-Vector, 60

P

P, 18
 P vs. NP, 67
 padding, 34
pairwise uniform, 78
 palindrome, 12, 16, 20, 37
partial functions, 17
 PCkt, 22
 PCP theorem, 62
 PH, 81
 PTime, 81
 PIT, 31
 polynomial identity testing, 31
 power hierarchy, 81
 PRG, 151
 PRGs from hard functions, 155
 prime number theorem, 32
 probabilistically-checkable-proofs, 62
 probability bounds for the deviation of the
 sum of random variables, 27
 pseudorandom, 151
pseudorandom generator, 151
 PSpace, 96

Q

quasi-linear time, 69
 Quasilinear-time completeness, 75

R

RAM, 24
random, 160
randomized TMs, 30
randomness, 26
Rapid-access machines, 24
reduction, 46
relations, 18
resamplable, 156
RRAM, 27

S

search problems, 61
Search-3Sat, 61
seed length, 151
SETH, 51
 Σ Time, 81
sorting, 71
sorting network, 72
Space, 96
Squaring, 48
stretch, 151
Strong exponential-time hypothesis, 51
structured objects, 17
SubquadraticTime, 49
Subset-sum, 53
Succinct-3Sat problem, 77
System, 59

T

tape machine, 15
TC, 122
tells, 151
terms (in a DNF), 22
 $\text{CktGates}(g(n))$, 22
Majority, 12
threshold circuit, 122
Time, 26
Time complexity, 18
time hierarchy, 40
TM, 15
TM-Time, 18
total functions, 17

Turing machine, Definition 2.1, 15

U

Unique-3Sat, 78
Unique-CktSat, 78
universal TM, 19

X

XOR Lemma, 160