

Computational Complexity Lecture Notes

Michael Levet

April 15, 2021

Contents

1	Combinatorial Circuits	3
1.1	Introduction to Circuits	3
1.2	Measures of Circuit Complexity	4
1.2.1	Exercises	5
1.3	Boolean Normal Forms	7
1.3.1	POSE and SOPE Normal Forms	7
1.3.2	Ring-Sum Expansion	9
1.3.3	Exercises	9
1.4	Parallel Prefix Circuits	11
1.4.1	Exercises	14
1.5	Parallel Prefix Addition	16
1.5.1	Exercises	17
1.6	Circuit Lower Bounds	18
1.7	Chapter Exercises	19
2	Computability	22
2.1	Turing Machines	22
2.1.1	Deterministic Turing Machine	22
2.1.2	Multitape Turing Machine	24
2.1.3	Non-deterministic Turing Machines	26
2.1.4	Exercises	27
2.2	Undecidability	28
2.3	Reducibility	29
2.3.1	Exercises	30
2.4	Oracle Turing Machines	32
2.4.1	Exercises	33
2.5	Arithmetic Hierarchy	34
2.5.1	Exercises	34
3	Structural Complexity	35
3.1	Ladner's Theorem	35
3.1.1	Exercises	37
3.2	Introduction to Space Complexity	38
3.2.1	PSPACE	39
3.2.2	L and NL	40
3.2.3	Exercises	41
3.3	Baker-Gill-Solovay Theorem	43
3.3.1	Exercises	44
3.4	Polynomial-Time Hierarchy: Introduction	45
3.4.1	Σ_i^p	46
3.4.2	Π_i^p	46
3.4.3	Exercises	47
3.5	Structure of Polynomial-Time Hierarchy	48
3.5.1	Exercises	49
3.6	Polynomial-Time Hierarchy and Oracles	50

3.6.1	Exercises	51
3.7	Time Hierarchy Theorem	52
3.7.1	Exercises	52
4	Interactive Proofs	54
4.1	Preliminaries	54
4.1.1	Exercises	55
4.2	Complexity Class IP	56
4.2.1	Exercises	56
4.3	IP = PSPACE	57
4.3.1	Exercises	60
4.4	Public Coins and Arthur-Merlin Protocols	61
4.4.1	Exercises	64
4.5	Arthur-Merlin Protocols and the Polynomial-Time Hierarchy	65
4.5.1	Exercises	65
5	Circuit Complexity: Razborov-Smolensky	66
5.1	Razborov-Smolensky: Introduction	66
5.2	Razborov-Smolensky: Approximating $AC^0[p]$ Circuits	68
5.2.1	Exercises	69
5.3	Razborov-Smolensky: Hilbert Functions	70
5.3.1	Exercises	72
5.4	Razborov-Smolensky: Obtaining Circuit Lower Bounds	73
5.4.1	Exercises	74
6	Circuit Complexity: The Power of Advice	75
6.1	Computation with Advice	75
6.1.1	Exercises	77
6.2	Sparse Sets and Mahaney's Theorem	78
6.2.1	Exercises	80
6.3	Karp-Lipton Theorems	81
6.3.1	Exercises	82

1 Combinatorial Circuits


These notes follow closely [Sav97][Chapter 2]. Both [Lev18, Ros12] also served as key references.

1.1 Introduction to Circuits

Definition 1. A *logic circuit* is a directed acyclic graph whose vertices are labeled with the names of Boolean functions or variables (inputs). The vertices labeled with the names of Boolean functions are referred to as *logic gates*.

Remark 2. The above definition does not clearly highlight the difference between a circuit and a logic gate. In practice, logic gates often compute simple functions, such as AND, OR, and NOT. The goal then is to design circuits using these elementary building blocks to compute more complicated Boolean functions.

Example 3. Below is an example of a Boolean circuit. The AND and OR gates are clearly labeled. The NOT gates are denoted by the white circles.



Our goal in this section is to establish that Boolean circuits are a universal model of computation. That is, a function is Turing computable if and only if it is computable by an equivalent family of Boolean circuits. To this end, we seek to construct a way to convert a Turing Machine to an equivalent family of circuits, as well as to construct an equivalent Turing Machine from a given family of circuits. We begin by trying to understand the circuit model of computation, and then proceed to examine the Turing Machine model. Our first step in this translation is to convert Boolean circuits to a procedural algorithmic construct, known as a straight-line program.

Definition 4. A *straight-line program* is a sequence of steps, each of which is of one of the following forms:

- Input Step: (s READ x), where x denotes an input variable,
- Output Step: (s OUTPUT i),
- Computation Step: (s OP i, \dots, k).

Here, OP is the given operation being performed. Now s is the index or line number of the given step, and i, \dots, k reference values computed at previous steps. So necessarily, $s \geq i, \dots, k$.

Example 5. Consider again the circuit from Example 3. We first examine a functional description of the given circuit.

$$\begin{aligned} g_1 &:= x \\ g_2 &:= y \\ g_3 &:= \bar{x} \\ g_4 &:= \bar{y} \\ g_5 &:= g_1 \wedge g_4 \\ g_6 &:= g_3 \wedge g_2 \\ g_7 &:= g_5 \vee g_6. \end{aligned}$$

In particular, the circuit first reads in the inputs and computes the relevant negations. The two AND gates are executed. Finally, the last OR gate is executed. We note that a circuit would actually execute g_1 and g_2 at the same time. Similarly, g_3 and g_4 would be executed in parallel, as would g_5 and g_6 . Note that a straight-line program, as well as an sequential model of computation like a RAM or Turing Machine, would not execute these statements in parallel. When equating two models of computation, the goal is to show that they compute the same family of functions. Two different models need not (and usually, do not) compute a given function in the same manner. We may also ask as to the cost, in time or space, in converting between two models. This is a question we will address later.

Next, we construct the equivalent straight-line program for this circuit.

```

(1 READ  $x$ )
(2 READ  $y$ )
(3 NOT 1)
(4 NOT 2)
(5 AND 1, 4)
(6 AND 3, 2)
(7 OR 5, 6)
(8 OUTPUT 7)

```

Remark 6. Informally, it may seem apparent at this point that every Boolean circuit computes a function of the form $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$. We will formalize this shortly. Perhaps less apparent is that for every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$, there exists a Boolean circuit that computes f . This second direction motivates several questions, including both how to construct such circuits and measures of circuit complexity.

We proceed first by showing that every Boolean circuit computes a function of the form $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$.

Definition 7. Let \mathcal{C} be a Boolean circuit. Let g_s be the function computed by the s th step of the straight-line program \mathcal{P} corresponding to \mathcal{C} . If the s th step is (s READ x), then $g_s = x$. If it is the computation step (s OP i, \dots, k), then $g_s = \text{OP}(g_i, \dots, g_k)$, where g_i, \dots, g_k are the functions computed at steps $i, \dots, k \leq s$.

If \mathcal{P} has n inputs and m outputs, then \mathcal{P} computes a function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$. If s_1, \dots, s_m are the output steps, then $f = (g_{s_1}, \dots, g_{s_m})$.

Remark 8. In light of this definition, we have that the function computed by the Boolean circuit \mathcal{C} is precisely the function computed by the straight-line program \mathcal{C} .

We now turn to showing that every Boolean function can be computed by a family of Boolean circuits. Any Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ can be expressed using a truth table. There are 2^n such rows in the truth table. One obvious approach is to build a circuit for each row, and then use AND gates on the outputs for the row circuits to generate the specified outputs for the function. Informally, such a circuit seems expensive, as we would require $\Theta(2^n)$ logic gates. We introduce some measures of circuit complexity to precisely measure the succinctness of a circuit.

1.2 Measures of Circuit Complexity

In order to measure the complexity of a circuit, we must first specify the allowed logic gates. This brings us to the notion of a basis.


Definition 9. The *basis* Ω of a circuit is the set of operations that it can use. Bases for Boolean circuits may only use Boolean functions. The *standard basis* is $\Omega_0 := \{\text{AND}, \text{OR}, \text{NOT}\}$. We say that a basis Ω is *complete* if every Boolean function can be realized using precisely the operations in Ω .

Once we have a basis, we may then begin to discuss measures of circuit complexity.

Definition 10. Let \mathcal{C} be a circuit. The *depth* of \mathcal{C} is the number of gates on the longest path through the circuit. The *circuit size* of \mathcal{C} is the number of gates it contains. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}^m$ be a Boolean function, and let Ω be a basis. The *circuit depth* of f with respect to the basis Ω , denoted $D_\Omega(f)$, is the minimum depth taken over all circuits that compute f . Similarly, the *circuit size* of f with respect to the basis Ω , denoted $C_\Omega(f)$, is the minimum size taken over all circuits that compute f .

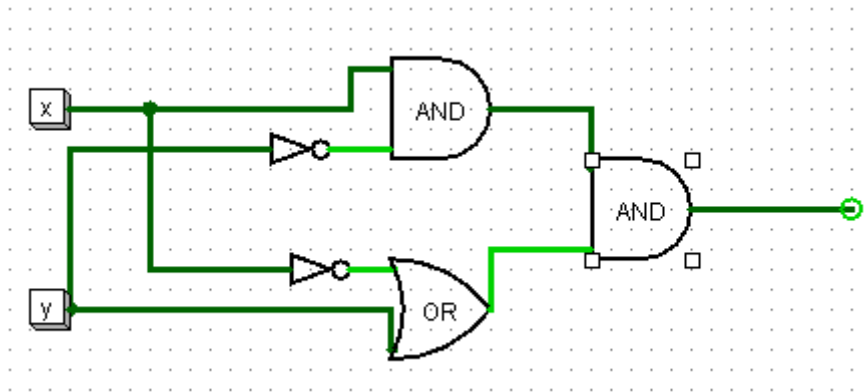
Remark 11. We note that in general, a circuit realizing $C_{\Omega}(f)$ may not also realize $D_{\Omega}(f)$.

Example 12. Again consider the circuit \mathcal{C} from Example 3. Here, the depth of \mathcal{C} is 3, realized by a sequence of a NOT gate, followed by an AND gate, and then an OR gate. The size of \mathcal{C} is 5, as there are two NOT gates, two AND gates, and a single OR gate.



1.2.1 Exercises

(Recommended) Problem 1. Consider the following circuit.



Do the following.

- Give a functional description of this circuit. (See Example 2 in the notes.)
- Give an equivalent straight-line program corresponding to this circuit.

(Recommended) Problem 2. Recall that the standard basis is $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$. Let \mathcal{F} be the set of Boolean functions with codomain $\{0, 1\}$ that are realizable over Ω_0 .¹

- Show that $f \in \mathcal{F}$ if and only if f is realizable over the basis $\Omega = \{\text{NAND}\}$.
- Let \mathcal{C} be a circuit realizing $C_{\Omega_0}(f)$, and let \mathcal{C}' be the circuit realizing f over the basis $\Omega = \{\text{NAND}\}$ corresponding to your transformation in the previous part. Relate the circuit size of \mathcal{C}' to $C_{\Omega_0}(f)$.
- Show that $f \in \mathcal{F}$ if and only if f is realizable over the basis $\Omega = \{\text{XOR}, \text{NOT}\}$.
- Let \mathcal{C} be a circuit realizing $C_{\Omega_0}(f)$, and let \mathcal{C}' be the circuit realizing f over the basis $\Omega = \{\text{XOR}, \text{NOT}\}$ corresponding to your transformation in the previous part. Relate the circuit size of \mathcal{C}' to $C_{\Omega_0}(f)$.

(Recommended) Problem 3. A *Half-Adder* is a circuit that adds to single-digit binary numbers x and y . The Half-Adder outputs the sum $x + y \pmod{2}$, as well as the carry bit. Note that the carry bit is a 1 precisely if a carry is generated by the addition .

- Implement the Half-Adder using only the AND, OR, and NOT gates.

¹This is, in fact, all such Boolean functions. However, we have not proven this yet.

- (b) Implement the Half-Adder, this time using the XOR gate. You may still use the AND, OR, and NOT gates. Comment on the difference in the number of gates required between your implementations in part (a) vs. part (b).

(Recommended) Problem 4. In this problem, we show the importance of having the NOT function in the standard basis. The *monotone basis* $\Omega_{\text{mon}} = \{\text{AND}, \text{OR}\}$, and we refer to a Boolean function using only the AND and OR operations as a *monotone Boolean function*. Do the following.

- (a) Prove that if $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a monotone Boolean function, then we may write:

$$f(x_1, x_2, \dots, x_n) = f(0, x_2, \dots, x_n) \vee (x_1 \wedge f(1, x_2, \dots, x_n)).$$

- (b) Denote \preceq to be the ordering on $\{0, 1\}^n$ where $x \preceq y$ if $x_i \leq y_i$ for all $i \in [n]$. Prove by induction on n that if $x, y \in \{0, 1\}^n$ satisfy $x \preceq y$; then for any monotone Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, we have that $f(x) \leq f(y)$.
- (c) Deduce that NOT is not a monotone Boolean function. Conclude that Ω_{mon} is not a complete basis.

1.3 Boolean Normal Forms

Our goal is to show that every Boolean function can be realized using a Boolean circuit. One strategy to answer this question is as follows. Does there exist a basis Ω , such that every Boolean function can be placed into some standard/normal form with respect to Ω ? A positive answer to this question would immediately allow us to construct a circuit for a given Boolean function. We introduce several normal forms, including the Disjunctive and Conjunctive Normal Forms, the Product of Sums Expansion, the Sum of Products Expansion, and the Ring-Sum Expansion.

1.3.1 POSE and SOPE Normal Forms

In this section, we introduce the Product of Sums Expansion (POSE) and Sum of Products Expansion (SOPE). We will also discuss the Conjunctive Normal Form (CNF), which is a special case of a POSE; as well as the Disjunctive Normal Form (DNF), which is a special type of SOPE. We begin with the definitions of POSE.

Definition 13 (Clause). A *clause* is a Boolean function $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ where φ consists of input variables or their negations, all added together (where addition is the OR operation).

Definition 14. Let $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. A *Product of Sums Expansion* (or *POSE*) of φ is the conjunction of the clauses C_1, \dots, C_k , such that :

$$\varphi = C_1 \wedge C_2 \wedge \dots \wedge C_k.$$

It is helpful to think of a POSE as an AND of ORs.

Example 15. The following functions are all in POSE form.

- $(A \vee \neg B \vee C) \wedge (\neg D \vee E \vee F)$
- x
- $x \vee y$
- $(x \vee y) \wedge z$

In contrast, the following functions are **not** in POSE form.

- $\neg(x \vee y)$ is not in CNF, as the OR is nested within the NOT. Note that $\neg(x \vee y)$ could be written in POSE form, as follows: $\neg x \vee \neg y$.
- $(x \wedge y) \vee z$

We next introduce Sum of Products Expansion.

Definition 16. A *term* is a Boolean function $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ where φ consists of input variables or their negations, all multiplied together (where multiplication is the AND operation).

Definition 17 (Sum of Products Expansion). Let $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. A *Sum of Products Expansion* (or *SOPE*) of φ is the disjunction of the terms D_1, \dots, D_k such that:

$$\varphi = D_1 \vee D_2 \dots \vee D_k.$$

It is helpful to think of a SOPE as an OR of ANDs.

We now consider some examples of Boolean expressions in SOPE form.

Example 18. The following formulas are in SOPE form.

- $(x \wedge y \wedge \neg z) \vee (\neg a \wedge b \wedge c)$
- $(x \wedge y) \vee z$

- $x \wedge y$
- x

The following formulas are not in SOPE form.

- $\neg(x \vee y)$ is not in SOPE, as the OR is nested inside the NOT. Note that $\neg x \wedge \neg y$ is also in SOPE form.
- $x \vee (y \wedge (c \vee d))$ is not in SOPE, as the OR is nested inside the AND.

We next discuss how to construct the CNF and DNF representations of a Boolean function. We begin by examining the truth table of the Boolean function. The rows that evaluate to 1 provide the basis for the DNF, keeping the values for each variable in their respective columns. In order to compute the CNF, we examine the rows that evaluate to 0 and take the negations of each value.

Example 19. Consider the function $\varphi : \{0, 1\}^3 \rightarrow \{0, 1\}$, given by the following truth table.

x	y	z	φ
0	0	0	0
0	0	1	1
0	1	0	1
0	1	1	0
1	0	0	1
1	0	1	0
1	1	0	0
1	1	1	1

We now compute the CNF and DNF.

- **CNF:** To compute the CNF, we examine the rows that evaluate to 0. We include rows 1, 4, 6, and 7. For each row, we invert the literals. So if $x = 0$ in the given row, we record x . If instead $x = 1$, we record \bar{x} . These rows contribute the following clauses.

- **Row 1:** $(x \vee y \vee z)$
- **Row 4:** $(x \vee \bar{y} \vee \bar{z})$
- **Row 6:** $(\bar{x} \vee y \vee \bar{z})$
- **Row 7:** $(\bar{x} \vee \bar{y} \vee z)$

So the CNF formulation of φ is:

$$\varphi = (x \vee y \vee z) \wedge (x \vee \bar{y} \vee \bar{z}) \wedge (\bar{x} \vee y \vee \bar{z}) \wedge (\bar{x} \vee \bar{y} \vee z).$$

- **DNF:** To compute the DNF, we examine the rows that evaluate to 1. We include rows 2, 3, 5, and 8. These rows contribute the following terms.

- **Row 2:** $(\bar{x} \wedge \bar{y} \wedge z)$
- **Row 3:** $(\bar{x} \wedge y \wedge \bar{z})$
- **Row 5:** $(x \wedge \bar{y} \wedge \bar{z})$
- **Row 8:** $(x \wedge y \wedge z)$

So the DNF formulation of φ is:

$$\varphi = (\bar{x} \wedge \bar{y} \wedge z) \vee (\bar{x} \wedge y \wedge \bar{z}) \vee (x \wedge \bar{y} \wedge \bar{z}) \vee (x \wedge y \wedge z).$$

Remark 20. We note that for the constant function $\varphi = 1$, $\text{CNF}(\varphi) = 1$ (i.e., the empty product). Similarly, for the constant function $\varphi = 0$, $\text{DNF}(\varphi) = 0$ (i.e., the empty sum).

Remark 21. As any Boolean function can be expressed in either CNF or DNF, it follows that any Boolean function can be computed using a Boolean circuit over the standard basis $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$. We record this observation with the following theorem.

Theorem 22. Every Boolean function can be realized using a Boolean circuit over the standard basis $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$.

Remark 23. We note that in practice, the CNF and DNF representations of a Boolean function may not be the most concise POSE and SOPE representations.

1.3.2 Ring-Sum Expansion

In this section, we introduce the Ring-Sum Expansion of a Boolean function f , which provides a way to express f over the basis $\{\text{XOR}, \text{AND}\}$. We begin with the definition of the Ring-Sum Expansion.

Definition 24. The *Ring-Sum Expansion* of a Boolean function f is the XOR (\oplus) of a constant and products of unnegated variables of f .

Remark 25. The set $\{0, 1\}$ together with the operations of addition (\oplus) and multiplication (\wedge) constitute the field \mathbb{F}_2 . As a field is a ring, this motivates the terminology *Ring-Sum Expansion*.

The Ring-Sum Expansion can be constructed starting from the DNF representation of a Boolean function. We note that for $x \in \{0, 1\}$, $\bar{x} = 1 \oplus x$. We first replace any OR (\vee) operator with XOR (\oplus). We next replace each negated variable \bar{x} in the DNF with $1 \oplus x$, and then apply the distribution law. If a term appears twice, we may cancel it out using commutativity of \oplus and the fact that $x \oplus x = 0$.

Example 26. Consider the Boolean function $\varphi(x_1, x_2, x_3) = (\bar{x}_1 \vee x_2) \wedge x_3$. We note that the DNF representation of φ is:

$$\varphi = \bar{x}_1 x_2 x_3 \vee \bar{x}_1 \bar{x}_2 x_3 \vee x_1 x_2 x_3.$$

For succinctness, we note that a product corresponds to the \wedge operator. So for instance, $\bar{x}_1 x_2 x_3 := \bar{x}_1 \wedge x_2 \wedge x_3$. We now construct the Ring-Sum Expansion from the DNF representation of φ :

$$\begin{aligned} \varphi &= [(1 \oplus x_1)x_2x_3] \oplus [(1 \oplus x_1)(1 \oplus x_2)x_3] \oplus x_1x_2x_3 \\ &= [x_2x_3 \oplus x_1x_2x_3] \oplus [x_3 \oplus x_1x_3 \oplus x_2x_3 \oplus x_1x_2x_3] \oplus x_1x_2x_3 \\ &= x_3 \oplus x_1x_3 \oplus x_1x_2x_3. \end{aligned}$$

1.3.3 Exercises

(Recommended) Problem 5. Suppose φ is given by the following truth table.

x	y	z	φ
0	0	0	0
0	0	1	0
0	1	0	0
0	1	1	0
1	0	0	1
1	0	1	0
1	1	0	1
1	1	1	1

- (a) Write φ in CNF.
- (b) Write φ in DNF.

(Recommended) Problem 6. Consider the following Boolean functions in DNF. For each function ϕ , determine if there is a sequence of inputs such that ϕ evaluates to 1 on those inputs.

- (a) $x \wedge \neg x$
- (b) $(x_1 \wedge \neg x_2 \wedge x_3 \wedge x_4) \vee (\neg x_1 \wedge \neg x_2 \wedge x_1 \wedge x_3)$

$$(c) (x_1 \wedge x_2 \wedge x_3) \vee (\neg x_1 \wedge \neg x_2 \wedge \neg x_3)$$

(Recommended) Problem 7. Suppose $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ is a Boolean function in DNF. Describe a polynomial-time algorithm to check whether φ has a satisfying instance. That is, establish that $\text{DNF-SAT} \in \text{P}$.

(Recommended) Problem 8. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. Denote $\text{CNF}(f)$ and $\text{DNF}(f)$ to be the CNF and DNF representations of f , respectively.

- (a) Show that $\text{CNF}(f)$ is unique. That is, suppose that φ, φ' are CNF realizations of f . Show that a given clause C belongs to φ if and only if C belongs to φ' .
- (b) Show that $\text{CNF}(f) = \overline{\text{DNF}(f)}$.

(Recommended) Problem 9. Let $f_{\oplus}^{(n)} : \{0, 1\}^n \rightarrow \{0, 1\}$ be the parity function. That is,

$$f_{\oplus}^{(n)}(x_1, \dots, x_n) = \begin{cases} 0 & : \sum_{i=1}^n x_i \equiv 0 \pmod{2}, \\ 1 & : \text{otherwise.} \end{cases}$$

Our goal is to show that the SOPE representing f has exponentially many terms.

- (a) Prove by induction on n that $f_{\oplus}^{(n)}$ has 2^{n-1} terms in the DNF representation.
- (b) Our goal now is to show that the DNF representation of $f_{\oplus}^{(n)}$ cannot be simplified. Show that each variable uniquely specifies a term in the DNF representation of $f_{\oplus}^{(n)}$.
- (c) Argue by contradiction that no term in the SOPE of $f_{\oplus}^{(n)}$ has fewer than n variables.
- (d) Conclude that the DNF representation of $f_{\oplus}^{(n)}$ cannot be simplified. Deduce that the SOPE has exponentially many terms.

Remark 27. A similar approach in Problem 9 may be used to show that the POSE expansion of $f_{\oplus}^{(n)}$ has exponentially many terms.

(Recommended) Problem 10. Let:

$$f_{\vee}^{(n)}(x_1, \dots, x_n) := \bigvee_{i=1}^n x_i.$$

Do the following.

- (a) Find an equivalent expression for $x \vee y$ using the basis $\{\text{AND}, \text{XOR}\}$.
- (b) Prove by induction that the Ring-Sum Expansion of $f_{\vee}^{(n)}$ contains every product term, except for the constant 1.

1.4 Parallel Prefix Circuits

The circuit model allows for more natural analysis and discussion of parallel computation, compared to sequential models such as the RAM or Turing Machine models. To this end, we ask about functions which can be effectively parallelized. Certain algebraic operations that satisfy the associative law constitute one such class of functions. The underlying set, coupled with the operation, is referred to as a *semi-group*. Precisely, we have the following definition.

Definition 28. A *semi-group* (S, \odot) consists of a set S together with an associative, binary operator $\odot : S \times S \rightarrow S$.

Example 29. We recall several examples of semigroups.

- (a) The set $S = \{0, 1\}$ together with the AND operator constitutes a semi-group.
- (b) The set $S = \{0, 1\}$ together with the OR operator constitutes a semi-group.
- (c) The natural numbers $\mathbb{N} = \{0, 1, 2, \dots\}$ forms a semi-group under addition.
- (d) Let $\{0, 1\}^*$ denote the set of all *finite* binary strings. The set $\{0, 1\}^*$ forms a semi-group under the operation of string concatenation.

Using the associativity of a semi-group (S, \odot) , we seek to compute a parallel prefix function, which returns the sequence running products for $x_1 \odot x_2 \odot \dots \odot x_n$.

Definition 30. Let (S, \odot) be a semi-group. A *prefix function* $\mathcal{P}_{\odot}^{(n)} : S^n \rightarrow S^n$ maps the input $(x_1, \dots, x_n) \mapsto (y_1, \dots, y_n)$; where for $1 \leq i \leq n$, $y_i = x_1 \odot x_2 \odot \dots \odot x_i$.

Suppose we have a semi-group (S, \odot) , as well as a logic gate for \odot with fan-in 2. Then in order to compute:

$$x_1 \odot x_2 \odot \dots \odot x_n,$$

we may use a binary tree where the leaves are the inputs and the root node is the output. Such a circuit has depth $\lceil \log_2(n) \rceil$. We require $\Theta(n^2)$ gates to return the sequence of running products. We leave the precise details in designing and analyzing such a circuit as an exercise for the reader.

We now ask whether we can use asymptotically fewer gates. Using a dynamic programming technique, we can construct an equivalent circuit with depth $2\lceil \log_2(n) \rceil$ that uses $\Theta(n)$ gates. While the depth of the circuit doubles, we note that the depth our new circuit is still $\Theta(\log(n))$. We now turn to constructing such a circuit. We begin with the definition of a parallel prefix circuit.

Our goal is to use dynamic programming to compute y_j . The circuit construction is recursive in nature. We introduce key observations and the high-level strategy before proceeding into designing the circuit. Define $x[r, r] = x_r$; and for $r \leq s$, let $x[r, s] := x_r \odot x_{r+1} \odot \dots \odot x_s$. So we have that $y_j := x[1, j]$. Now as \odot is associative, we have that $x[r, s] = x[r, t] \odot x[t+1, s]$, for $r \leq t < s$. To make this clear, we have by definition that:

$$\begin{aligned} x[r, s] &= x_r \odot x_{r+1} \odot \dots \odot x_s, \\ x[r, t] &= x_r \odot x_{r+1} \odot \dots \odot x_t, \\ x[t+1, s] &= x_{t+1} \odot x_{t+2} \odot \dots \odot x_s. \end{aligned}$$

So:

$$\begin{aligned} x[r, s] &= \left(x_r \odot x_{r+1} \odot \dots \odot x_t \right) \odot \left(x_{t+1} \odot x_{t+2} \odot \dots \odot x_s \right) \\ &= x[r, t] \odot x[t+1, s]. \end{aligned}$$

We now turn to designing the circuit itself. Note that we assume the number of inputs n to be a power of 2; that is, $n = 2^k$. If the number of inputs is not a power of 2, we may add additional inputs until we have 2^k inputs for some k . When examining the outputs of the prefix function, we need only consider y_n ; we may

ignore y_i for $i > n$.

The construction of our circuit to compute $\mathcal{P}_{\odot}^{(n)}$ is recursive in nature. Suppose we have inputs x_1, \dots, x_n . We obtain the following $(n/2)$ -tuple:

$$(x[1, 2], x[3, 4], \dots, x[2^k - 1, 2^k]).$$

Now observe that $x[i, i + 1] = x[i, i] \odot x[i + 1, i + 1]$. So suppose we compute $\mathcal{P}_{\odot}^{(n/2)}(x[1, 2], \dots, x[2^k - 1, 2^k])$. Recall that the output of $\mathcal{P}_{\odot}^{(n/2)}(x[1, 2], \dots, x[2^k - 1, 2^k]) = (z_1, \dots, z_{n/2})$, where:

$$\begin{aligned} z_i &= x[1, 2] \odot x[3, 4] \odot \dots \odot x[2i - 1, 2i] \\ &= x[1, 2i]. \end{aligned}$$

That is,

$$\mathcal{P}_{\odot}^{(n/2)}(x[1, 2], \dots, x[2^k - 1, 2^k]) = (x[1, 2], x[1, 4], \dots, x[1, 2^k]).$$

Note that $\mathcal{P}_{\odot}^{(n)}$ must also compute $x[1, 1], x[1, 3], x[1, 5], \dots, x[1, 2^k - 1]$. Thankfully, these are easy to compute once we have computed $x[1, 2], x[1, 4], \dots, x[1, 2^k]$. Namely, observe that $x[1, 2i + 1] = x[1, 2i] \odot x[2i + 1]$. Note that as $x[1, 1] = x_1$, we do not require an operation to compute $x[1, 1]$. Thus, only $2^{k-1} - 1$ operations are required to compute $x[1, 1], x[1, 3], x[1, 5], \dots, x[1, 2^k - 1]$.

Before analyzing the circuit's complexity, we consider an example.

Example 31. Suppose we have $n = 2^3$ inputs, x_1, \dots, x_8 . Our goal is to compute $\mathcal{P}_{\odot}^{(8)}(x_1, \dots, x_8)$. The circuit proceeds as follows.

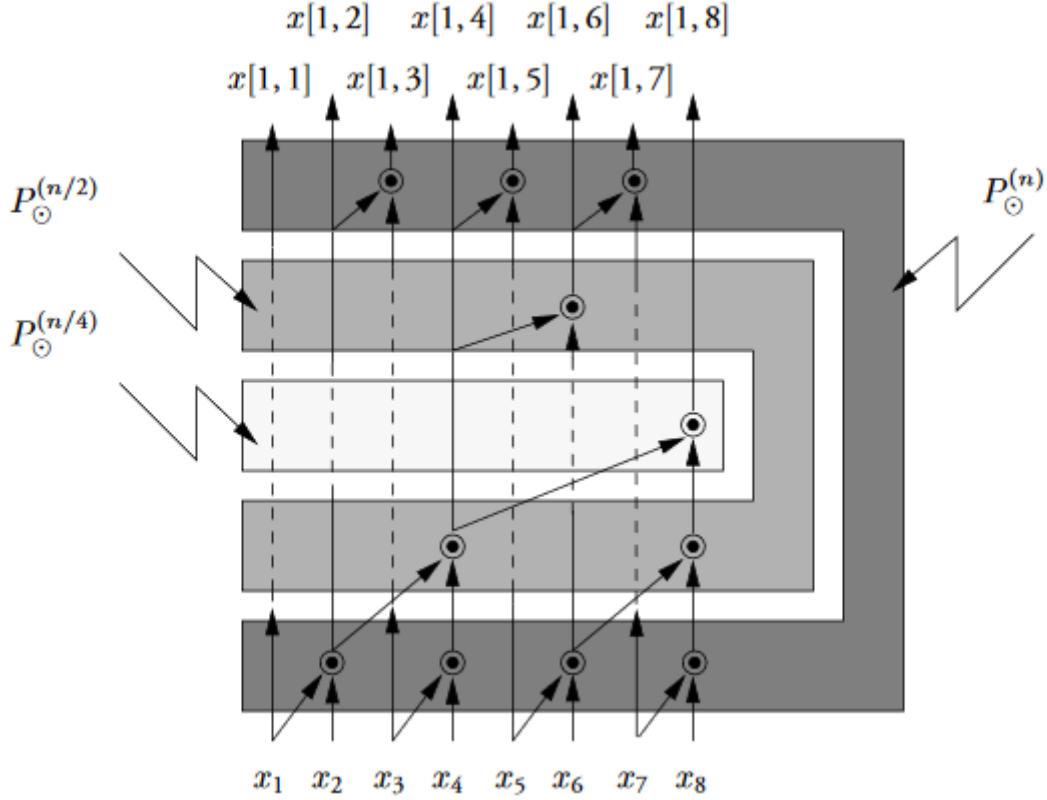
- We first compute $x[1, 2], x[3, 4], x[5, 6], x[7, 8]$. In order to more easily follow the work in the recursive calls, we label $a_1 = x[1, 2], a_2 = x[3, 4], a_3 = x[5, 6], a_4 = x[7, 8]$.
- The $\mathcal{P}_{\odot}^{(8)}(x_1, \dots, x_8)$ circuit now recursively computes $\mathcal{P}_{\odot}^{(4)}(a_1, a_2, a_3, a_4)$. Our recursive call starts by computing:

$$\begin{aligned} a_1 \odot a_2 &\text{ (which we recall is } x[1, 2] \odot x[3, 4] = x[1, 4]) \\ a_3 \odot a_4 &\text{ (which we recall is } x[5, 6] \odot x[7, 8] = x[5, 8]) \end{aligned}$$

In order to more easily follow the recursive calls, we label $b_1 = x[1, 4]$ and $b_2 = x[5, 8]$.

- The $\mathcal{P}_{\odot}^{(4)}(a_1, a_2, a_3, a_4)$ circuit now recursively computes $\mathcal{P}_{\odot}^{(2)}(b_1, b_2) = (b_1, b_1 \odot b_2) = (x[1, 4], x[1, 8])$.
- We now return control to the $\mathcal{P}_{\odot}^{(4)}(a_1, a_2, a_3, a_4)$ circuit. We have computed $a[1, 2]$ and $a[1, 4]$. We note that $a[1, 1] = 1$, and so we only require one additional operation to compute $a[1, 3] = a[1, 2] \odot a_3$. So the $\mathcal{P}_{\odot}^{(4)}(a_1, a_2, a_3, a_4)$ circuit has completed its execution and returns control to $\mathcal{P}_{\odot}^{(8)}(x_1, \dots, x_8)$. Note that we have computed $a[1, 1] = x[1, 2]$, $a[1, 2] = x[1, 4]$, $a[1, 3] = x[1, 6]$, and $a[1, 4] = x[1, 8]$.
- Now that control has been returned to $\mathcal{P}_{\odot}^{(8)}(x_1, \dots, x_8)$, it remains to compute $x[1, 1], x[1, 3], x[1, 5]$, and $x[1, 7]$. We again note that $x[1, 1] = 1$, and so does not require any operations to compute. We need one operation to compute $x[1, 3] = x[1, 2] \odot x_3$, one operation to compute $x[1, 5] = x[1, 4] \odot x_5$, and one operation to compute $x[1, 7] = x[1, 6] \odot x_7$.
- $\mathcal{P}_{\odot}^{(8)}(x_1, \dots, x_8)$ now outputs $(x[1, 1], x[1, 2], x[1, 3], \dots, x[1, 8])$.

Pictorially, we include the diagram for this circuit from [Sav97][Chapter 2, Figure 2.13].



We now turn to analyzing the circuit construction.

- **Circuit Size:** We note that $\mathcal{P}_{\odot}^{(n)}$ uses an initial 2^{k-1} \odot gates to compute $x[1,2], x[3,4], \dots, x[2^k-1, 2^k]$. At the last layer, the circuit uses an additional $2^{k-1} - 1$ gates to compute $x[1,3], \dots, x[1, 2^k-1]$. No gate is needed to compute $x[1,1] = x_1$. This yields $2^k - 1$ total gates, in addition to the number of gates that the circuit for $\mathcal{P}_{\odot}^{(n/2)}$ uses. Therefore, the size of the circuit is given by the recurrence:

$$C(k) = \begin{cases} C(k-1) + 2^k - 1 & : k > 0, \\ 0 & : k = 0. \end{cases}$$

We stress that the input variable k for the recurrence $C(k)$ is the same k as in the exponent of $n = 2^k$. Solving this recurrence using the unrolling method yields the closed form expression: $C(k) = 2^{k+1} - k - 2$. Noting that $n = 2^k$ (and so, $k = \log_2(n)$), we have that:

$$\begin{aligned} 2^{k+1} - k - 2 &= 2 \cdot 2^k - k - 2 \\ &= 2n - \log_2(n) - 2. \end{aligned}$$

It follows that if Ω is a basis containing \odot , then $C_{\Omega}(\mathcal{P}_{\odot}^{(n)}) \leq 2n - \log_2(n) - 2$. We make no claims that the prefix circuit construction above is optimal. As a result, we only obtain an upper bound on $C_{\Omega}(\mathcal{P}_{\odot}^{(n)})$.

- **Circuit Depth:** We note that $\mathcal{P}_{\odot}^{(n)}$ has three layers:
 - The initial layer that computes $x[1,2], x[3,4], \dots, x[2^k-1, 2^k]$,
 - The layer computing $\mathcal{P}_{\odot}^{(n/2)}(x[1,2], x[3,4], \dots, x[2^k-1, 2^k])$, and
 - The layer computing $x[1,1], x[1,3], \dots, x[1, 2^k-1]$.

There are two layers, each of depth 1, not associated with the recursive call to $\mathcal{P}_{\odot}^{(n/2)}$. So the circuit depth satisfies the following recurrence:

$$D(k) = \begin{cases} D(k-1) + 2 & : k > 0, \\ 0 & : k = 0. \end{cases}$$

Again, we stress that the input variable k for the recurrence $C(k)$ is the same k as in the exponent of $n = 2^k$. Solving this recurrence using the unrolling method yields the closed form expression: $D(k) = 2k$. Noting that $k = \log_2(n)$, we have that $2k = 2\log_2(n)$. So the depth of our circuit is $2\log_2(n)$. It follows that if Ω is a basis containing \odot , then $D_\Omega(\mathcal{P}_\odot^{(n)}) \leq 2\log_2(n)$. We make no claims that the prefix circuit construction above is optimal. As a result, we only obtain an upper bound on $D_\Omega(\mathcal{P}_\odot^{(n)})$.

1.4.1 Exercises

(Recommended) Problem 11. Suppose we have the basis Ω with the associative, binary operator \odot . Suppose that the \odot gates have fan-in 2 (that is, the \odot gates accept two inputs). Do the following.

- Design a circuit of depth $\lceil \log_2(n) \rceil$ to compute $x_1 \odot x_2 \cdots \odot x_n$.
- Comment on why the associativity of \odot is a necessary assumption for your circuit's correctness in part (a).
- Suppose that $n = 2^k$ for some $k \in \mathbb{N}$. Determine the size of your circuit.
- Suppose now that n is not a power of 2. Determine a suitable function $f(n)$ such that the size of your circuit is $\Theta(f(n))$. **[Hint:** We note that there exists $k \in \mathbb{N}$ such that $2^k < n < 2^{k+1}$.]
- Now suppose instead that the \odot gate has unbounded fan-in. Design a circuit to compute $x_1 \odot x_2 \cdots \odot x_n$. What are the size and depth of your new circuit, and how do they compare to the size and depth of the circuit you designed in part (a)?

Definition 32. Let $k \in \mathbb{N}$. We say that a language $L \in \text{NC}^k$ if there exists a uniform family of circuits $(C_n)_{n \in \mathbb{N}}$ over the standard basis $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$ such that the following conditions hold:

- A string $\omega \in \{0, 1\}^*$ of length n is in L if and only if $C_n(\omega) = 1$ (that is, C_n on input ω evaluates to 1).
- Each circuit has fan-in 2 for the AND and OR gates, and fan-in 1 for the NOT gates.
- The circuit C_n has depth $O(\log^k(n))$ and size $O(n^m)$. The implicit constants and the exponent m both depend on L .

Definition 33. Let $k \in \mathbb{N}$. We say that a language $L \in \text{AC}^k$ if there exists a uniform family of circuits $(C_n)_{n \in \mathbb{N}}$ over the standard basis $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$ such that the following conditions hold:

- A string $\omega \in \{0, 1\}^*$ of length n is in L if and only if $C_n(\omega) = 1$ (that is, C_n on input ω evaluates to 1).
- Each circuit has fan-in 1 for the NOT gates.
- The AND and OR gates have unbounded fan-in.
- The circuit C_n has depth $O(\log^k(n))$ and size $O(n^m)$. The implicit constants and the exponent m both depend on L .

Remark 34. Note that the only difference between an NC circuit and an AC circuit is that the NC circuits have bounded fan-in for the AND and OR gates, while the AC circuits allow for unbounded fan-in for the AND and OR gates.

Remark 35. We also note that the uniformity condition means that there is a Turing Machine M such that on input 1^n , M can generate the circuit C_n . This condition is not important at present; however, it will become important later when we relate Turing Machines to circuits.

(Recommended) Problem 12. Fix $k \in \mathbb{N}$. Do the following.

- Show that $\text{NC}^k \subseteq \text{AC}^k$.
- Show that $\text{AC}^k \subseteq \text{NC}^{k+1}$.

Remark 36. Define:

$$\text{NC} := \bigcup_{k \in \mathbb{N}} \text{NC}^k.$$

In light of Problem 12, we have that:

$$\text{NC} = \bigcup_{k \in \mathbb{N}} \text{AC}^k.$$

Remark 37. It is known that $\text{NC}^0 \subsetneq \text{AC}^0 \subsetneq \text{NC}^1$. For $k \geq 1$, none of the containments $\text{NC}^k \subseteq \text{AC}^k \subseteq \text{NC}^{k+1}$ are known to be strict.

1.5 Parallel Prefix Addition

In this section, we examine a parallel prefix circuit for binary addition. Let $u \in \mathbb{N}$. We write the binary expansion of u as:

$$u = \sum_{i=0}^{n-1} u_i 2^i,$$

where each $u_i \in \{0, 1\}$ and $n := \lceil \log_2(u) \rceil$. Now let $v \in \mathbb{N}$, and consider the binary expansion of v :

$$v = \sum_{i=0}^{m-1} v_i 2^i.$$

For the rest of this section, we assume without loss of generality that $m = n$. Otherwise, we consider $\max\{m, n\}$. Our goal now is to obtain the binary representation of $u + v$, we denote as:

$$u + v = \sum_{i=0}^n s_i 2^i.$$

Our primary goal is to determine the values of the s_i terms. We note that when we evaluate $u_i + v_i$, a carry bit is generated for our evaluation of $u_{i+1} + v_{i+1}$. If $u_i = 1$ and $v_i = 1$, then a carry bit of 1 is generated. We denote this carry bit as c_{i+1} , as it will be factored into the calculation of s_{j+1} . If $i \geq 1$, then we also need to consider the carry bit generated at the $i - 1$ stage, which is c_i . If $c_i = 1$ and at least one of $u_i = 1$ or $v_i = 1$, then a carry bit of 1 is generated at stage i . That is, $c_{i+1} = 1$. Otherwise, $c_{i+1} = 0$. Effectively, our circuit will have to track both the carry bits and the s_i bits.

In order to compute these values, we introduce intermediary variables called *generate* and *propagate* bits.

- The i th generate bit is given by $g_i := u_i \wedge v_i$. That is, $g_i = 1$ if and only if adding u_i and v_i generates a 1 for the carry bit c_{i+1} .
- The i th propagate bit is given by $p_i := u_i \oplus v_i$. Observe that p_i and g_i cannot both be 1. So if $p_i = 1$, we have that either $u_i = 1$ or $v_i = 1$, but u_i and v_i are not both 1. It follows that the carry bit c_i from the previous stage determines whether $c_{i+1} = 1$.

With the propagate and generate bits defined, we observe the following.

- First, $s_i = p_i \oplus c_i$. We note that $p_i = u_i \oplus v_i$, so $s_i = u_i \oplus v_i \oplus c_i$. That is, we add u_i and v_i , and then add in the incoming carry bit.
- Second, we observe that the carry bit generated satisfies the following:

$$c_{i+1} = (p_i \wedge c_i) \vee g_i.$$

If $g_i = 1$, then both $u_i = 1$ and $v_i = 1$. So adding u_i and v_i generates a carry bit. Otherwise, $c_{i+1} = 1$ precisely if $u_i \oplus v_i = 1$ and $c_i = 1$.

We now turn to defining our parallel-prefix addition circuit. Rather than tracking the carry and s_j bits at each intermediary step, we instead track the propagate and generate bits at each stage. Precisely, we examine ordered pairs (p_i, g_i) , where again p_i is the i th propagate bit and g_i is the i th generate bit. Our goal is to design a circuit that outputs whether a carry bit was generated at each stage. Note that the propagate and generate bits p_i and g_i depend only on the u_i and v_i , the i th positions of the binary numbers we are adding. So we may easily compute p_i and g_i in parallel.

Suppose we are considering consecutive indices (p_i, g_i) and (p_{i+1}, g_{i+1}) . Note that $p_i = g_i$ precisely if $p_i = 0$ and $g_i = 0$. Similarly, $p_{i+1} = g_{i+1}$ precisely if $p_i = 0$ and $g_i = 0$. We may track whether a propagate was generated at stages i and $i + 1$ by considering $p_i \wedge p_{i+1}$. Now to check whether the carry bit c_{i+2} takes on the value 0 or 1, when restricting attention to positions i and $i + 1$, we check that $g_{i+1} = 1$ (which means that $u_i \wedge v_i = 1$) or $p_{i+1} \wedge g_i = 1$. Note that if $g_i = 1$, then $c_{i+1} = 1$. So $p_{i+1} \wedge g_i = 1$ implies that $p_{i+1} \wedge c_{i+1} = 1$.

So the $(i + 1)$ st carry bit propogates.

In order to define a parallel-prefix circuit, we first need a suitable associative operator. Define $\diamond : \{0, 1\}^2 \times \{0, 1\}^2 \rightarrow \{0, 1\}^2$ by:

$$(a, b) \diamond (c, d) = (a \wedge c, (b \wedge c) \vee d).$$

Again, we write $a \wedge c$ as ac , to improve readability. We check that \diamond is associative. Observe that:

$$\begin{aligned} ((a, b) \diamond (c, d)) \diamond (e, f) &= (a, b) \diamond ((c, d) \diamond (e, f)) \\ &= (ace, bce \vee de \vee f). \end{aligned}$$

We now define our lookup table. Denote $\pi[j, j] = (p_j, g_j)$. For $j < k$, let $\pi[j, k] = \pi[j, k - 1] \diamond \pi[k, k]$. By induction, we have the following:

- The first component of $\pi[j, k]$ is 1 if and only if a carry propogates through each stage $j, j + 1, \dots, k$.
- The second component of $\pi[j, k]$ is 1 if and only if a carry is generated at some stage r , where $j \leq r \leq k$; and that carry propogates from stage r through stage k .

We apply the parallel-prefix circuit from Section 1.4 with the associative operator \diamond . The circuit takes as input $(\pi[0, 0], \pi[1, 1], \dots, \pi[n - 1, n - 1])$ and returns the sequence $(\pi[0, 0], \pi[0, 1], \dots, \pi[0, n - 1])$. We note that there is no carry bit of 1 when adding the ones place bits u_0 and v_0 . So $c_0 = 0$. As the first component of $\pi[0, i]$ tracks whether a carry bit value of 1 propogated through each stage $0, \dots, i$, we have that the first component of $\pi[0, i]$ is 0.

The second component of $\pi[0, i]$ tracks whether a a carry bit is generated on or before stage i . By unrolling the expression for c_{j+1} :

$$\begin{aligned} c_{j+1} &= p_j c_j \vee g_j \\ &= p_j (p_{j-1} c_{j-1} \vee g_{j-1}) \vee g_j = p_j p_{j-1} c_{j-1} \vee g_j, \end{aligned}$$

we see that the expression in the second component of $\pi[0, i]$ is precisely c_{j+1} . This unrolling also highlights the necessity of recording whether a carry propogates through each stage in the first component. Now the bit s_i may be recovered from taking the Exclusive Or of p_i and the second component of $\pi[0, i - 1]$.

1.5.1 Exercises

(Recommended) Problem 13. Recall the associative operator $\diamond : \{0, 1\}^2 \times \{0, 1\}^2 \rightarrow \{0, 1\}^2$, given by:

$$(a, b) \diamond (c, d) = (ac, bc \vee d).$$

Denote $\pi[j, j] = (p_j, g_j)$. For $j < k$, let $\pi[j, k] = \pi[j, k - 1] \diamond \pi[k, k]$.

- Show by induction on $\ell := k - j$ that the first component of $\pi[j, k]$ is 1 precisely if a carry bit of 1 propogates through the full adder stages indexed j, \dots, k .
- Show by induction on $\ell := k - j$ that the second component of $\pi[j, k]$ is 1 precisely if a carry bit of 1 is generated at some stage r , where $j \leq r \leq k$, and that carry propogates from stage r through stage k .

1.6 Circuit Lower Bounds

In this section, we show that most Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ require size exponential in n . The key idea is that there are more Boolean functions than there are small circuits. We note that there are 2^{2^n} such Boolean functions. We modify the proof from Kopparty, Kim, and Lund [KKL13], which is considerably simpler than that found in Savage [Sav97][Theorem 2.12.1].

We begin by analyzing the circuit size.

Theorem 38 (Shannon's Theorem). Let $\epsilon \in (0, 1)$. The fraction of Boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ that have size complexity satisfying:

$$C_{\Omega_0}(f) > \frac{2^n}{2n}(1 - \epsilon)$$

is at least $1 - 2^{-\epsilon \cdot 2^n}$ when $n \geq 6$.

Proof. Let $G \geq 0$, and denote $M(G)$ to be the number of Boolean functions that are computable using a circuit of at most G gates. Each gate accepts at most 2 inputs, and the order in which the inputs are selected does not matter. Now an input to a given gate may come from either another gate or a wire corresponding to a variable. There are $n + G - 1$ such selections. So there are at most $\binom{n+G-1}{2} \leq \binom{n+G}{2}$ ways to select the inputs for a given gate. Now there are three gates in the standard basis Ω_0 . So there are at most $3\binom{n+G}{2}$ ways to choose gates.

Now take:

$$G = \frac{2^n}{2n}(1 - \epsilon).$$

We show the number of circuits of size at most G is asymptotically smaller than the total number of Boolean functions on n variables. We note that there are 2^{2^n} Boolean functions on n variables. Now we have that:

$$M(G) \leq 3^G \binom{n+G}{2}^G \tag{1}$$

$$< 3^G (n+G)^{2G} \tag{2}$$

$$< 3^G (2G)^{2G} \tag{3}$$

$$< 3^{2G} (2G)^{2G} \tag{4}$$

$$= 6^{2G} \cdot G^{2G} \tag{5}$$

$$= 6^{2 \cdot 2^n(1-\epsilon)/(2n)} \cdot \left(\frac{2^n(1-\epsilon)}{n} \right)^{2 \cdot 2^n(1-\epsilon)/(2n)} \tag{6}$$

$$= 6^{2^n(1-\epsilon)/n} \cdot \left(\frac{2^n(1-\epsilon)}{n} \right)^{2^n(1-\epsilon)/n} \tag{7}$$

$$= \left(\frac{6(1-\epsilon)}{n} \right)^{2^n(1-\epsilon)/n} \cdot (2^n)^{2^n(1-\epsilon)/n} \tag{8}$$

$$= \left(\frac{6(1-\epsilon)}{n} \right)^{2^n(1-\epsilon)/n} \cdot 2^{2^n(1-\epsilon)} \tag{9}$$

$$\leq 2^{2^n(1-\epsilon)}. \tag{10}$$

We note that the upper bound at line (10) holds whenever $n \geq 6$. As $M(G) < 2^{2^n(1-\epsilon)}$, it follows that the fraction of Boolean functions requiring more than G gates is:

$$1 - 2^{-\epsilon \cdot 2^n},$$

as desired. □

1.7 Chapter Exercises

(Recommended) Problem 14. It is often desirable to take a circuit over the standard basis $\{\text{AND}, \text{OR}, \text{NOT}\}$ and convert it to an equivalent circuit without the negations. This motivates *dual-rail logic* circuits, which allow us to effectively push the negations down to the inputs. In dual-rail logic, the variable $|x\rangle$ is represented by the pair (x, \bar{x}) . In particular, observe that $|0\rangle = (0, 1)$ and $|1\rangle = (1, 0)$. We now turn to defining the dual-rail logical operations.

- The DRL-AND operation \wedge is defined by $|x\rangle \wedge |y\rangle = |x \wedge y\rangle$. Note that if $|x\rangle = (x, \bar{x})$ and $|y\rangle = (y, \bar{y})$, then we may realize $|x \wedge y\rangle$ over a classical circuit, using the standard basis $\Omega_0 := \{\text{AND}, \text{OR}, \text{NOT}\}$ as follows:

$$|x\rangle \wedge |y\rangle = (x \wedge y, \overline{x \wedge y}) = (x \wedge y, \bar{x} \vee \bar{y}).$$

- The DRL-OR operation \vee is defined by $|x\rangle \vee |y\rangle = |x \vee y\rangle$.
- The DRL-NOT operation can be realized by physically twisting the wires on (x, \bar{x}) , and so we do not include DRL-NOT in our basis.

Do the following.

- Show how to realize DRL-OR over a classical circuit, using the standard basis $\{\text{AND}, \text{OR}, \text{NOT}\}$.
- Deduce that every Boolean function $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}^m$ can be realized by a dual-rail logic circuit in which the standard NOT gates are only used on input variables (to obtain the pair (x, \bar{x})).
- Let $\text{DRL} = \{\text{DRL-AND}, \text{DRL-OR}\}$ be the standard dual-rail logic basis. Let $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$. Show that: $C_{\text{DRL}}(\varphi) \leq 2C_{\Omega_0}(\varphi)$.
- Show that $D_{\text{DRL}}(\varphi) \leq D_{\Omega_0}(\varphi) + 1$. For circuits over Ω_0 , do not count the NOT gates as contributing to the depth.²

(Recommended) Problem 15. For this problem, we restrict attention to the basis $\Omega = \{\text{AND}, \text{OR}, \text{NOT}, \text{XOR}\}$, where the AND, OR, and XOR gates all have fan-in 2. We are interested in the Membership problem, which takes as input $x \in \{0, 1\}^n$ and $y \in \{0, 1\}$; the goal is to decide whether there exists an $i \in [n]$ such that $x_i = y$.

- Design a circuit that computes the function $f_{\text{equality}}^{(n)} : \{0, 1\}^n \times \{0, 1\} \rightarrow \{0, 1\}$, where $f_{\text{equality}}(x, y) = 1$ if and only if $x = y$. [**Hint:** Start with the case when $n = 1$.]
- The *membership* function $f_{\text{membership}}^{(n)} : \{0, 1\}^n \times \{0, 1\} \rightarrow \{0, 1\}$ is given by $f_{\text{membership}}^{(n)}(x_1, \dots, x_n; y) = 1$ if and only if there exists an $i \in [n]$ such that $x_i = y$. Using part (a), design a circuit to compute $f_{\text{membership}}^{(n)}$.
- Analyze the size and depth of your circuit in part (b). In particular, conclude that the Membership problem is in NC^1 . [**Note:** NC^1 circuits are defined over the standard basis $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$. Why does allowing for XOR gates not change the result?]
- Strengthen the above result to show that the Membership problem is in AC^0 . That is, if we allow for unbounded fan-in, then we only require a constant depth circuit. [**Note:** The same note from part (c) about the XOR gate applies to AC circuits as well.]

Definition 39. A function $\varphi(x_1, \dots, x_n)$ is said to be *symmetric* if for every permutation $\pi \in \text{Sym}(n)$, we have that:

$$\varphi(x_1, \dots, x_n) = \varphi(x_{\pi(1)}, \dots, x_{\pi(n)}).$$

The building block for our symmetric functions are the *elementary symmetric functions* $e_t^{(n)} : \{0, 1\}^n \rightarrow \{0, 1\}$, where $0 \leq t \leq n$, given by:

$$e_t^{(n)}(x_1, \dots, x_n) = \begin{cases} 1 & \text{if } \sum_{i=1}^n x_i = t, \\ 0 & \text{otherwise.} \end{cases}$$

²This is a standard assumption in Circuit Complexity.

The next problems serve as practice for using elementary symmetric functions to design circuits that compute more complicated symmetric functions.

(Recommended) Problem 16. Let $n \in \mathbb{Z}^+$, and let $1 \leq t \leq n$. The *threshold function* is given by:

$$\tau_t^{(n)}(x_1, \dots, x_n) = \begin{cases} 1 & : \text{if } \sum_{i=1}^n x_i \geq t, \\ 0 & : \text{otherwise.} \end{cases}$$

Describe how to build an AC circuit using the standard basis and elementary symmetric functions.

(Recommended) Problem 17. Let $n \in \mathbb{Z}^+$. The *binary sort* function $f_{\text{sort}}^{(n)} : \{0, 1\}^n \rightarrow \{0, 1\}$ sorts an n -tuple into descending order. Here, for $x \in \{0, 1\}^n$, we have:

$$f_{\text{sort}}^{(n)}(x) = (\tau_1^{(n)}(x), \tau_2^{(n)}(x), \dots, \tau_n^{(n)}(x)).$$

Show that beyond the cost of implementing the elementary symmetric functions, we can realize $f_{\text{sort}}^{(n)}$ using an additional depth of $O(\log(n))$ and an additional $O(n)$ gates from the standard basis $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$.

(Recommended) Problem 18. Let $m, n \in \mathbb{Z}^+$, and let $0 \leq c \leq m - 1$. Define the *modulus function* to be:

$$f_{c, \text{ mod } m}^{(n)}(x_1, \dots, x_n) = \begin{cases} 1 & : \text{if } \sum_{i=1}^n x_i \equiv c \pmod{m}, \\ 0 & : \text{otherwise.} \end{cases}$$

Describe how to build a circuit using the standard basis and elementary symmetric functions. Pay close attention to the fact that n is *fixed*.

(Recommended) Problem 19. Recall from Problem 12 that $\text{NC}^0 \subseteq \text{AC}^0$. In this problem, we show that the containment is strict.

- (a) Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function. Show that if f is computed by an NC^0 circuit of depth d , then the circuit depends on at most 2^d inputs.
- (b) Show that the following function is not in NC^0 .

$$f_{\wedge}^{(n)}(x_1, \dots, x_n) = \bigwedge_{i=1}^n x_i.$$

- (c) Deduce that $\text{NC}^0 \subsetneq \text{AC}^0$.

(Recommended) Problem 20. In this problem, we seek to better understand the space between NC^1 and NC^2 . Let L be the set of languages decidable by deterministic Turing Machines³ where the input is read-only and only $O(\log(n))$ additional space is available for read and write access. Let NL denote the set of languages decidable by non-deterministic Turing Machines where the input is read-only and only $O(\log(n))$ additional space is available for read and write access. Equivocally, NL is the set of languages L ; where if $x \in L$, we can verify this in space $O(\log(n))$. Observe that:

$$\text{L} \subseteq \text{NL}.$$

Our goal will be to establish upper bounds for NL . To do this, we consider the **Connectivity** problem, which takes as input a directed graph $G(V, E)$ and two vertices $u, v \in V(G)$; we seek to decide whether there is a $u \rightarrow v$ path in G . The **Connectivity** problem is NL -complete under logspace reductions.

- (a) We think of a directed graph $G(V, E)$ as a binary relation $E \subseteq V(G) \times V(G)$. The *transitive closure* of E (which we may also refer to as the transitive closure of G) is the smallest transitive relation that contains E . Let $u, v \in V(G)$. Show that there is a $u \rightarrow v$ path if and only if (u, v) is in the transitive closure of E .

³You may think of Turing Machines as algorithms. We are not concerned with the technicalities of Turing Machines.

- (b) For the remainder of this problem, we will consider Boolean matrix multiplication, where addition is given by the OR operator and multiplication is given by the AND operator. Let A be the adjacency matrix of the graph G . Let $M := (A \vee I)$, where I is the identity matrix. Prove by induction that there is a directed path of length at most k from $u \rightarrow v$ if and only if $M_{uv}^k = 1$. It follows immediately that we can recover the transitive closure of G from M^n , where n is the number of vertices.
- (c) Let A, B be $n \times n$ Boolean matrices. Show that computing AB is in NC^1 . [**Hint:** When multiplying non-Boolean matrices X and Y , we have that $XY_{ij} = \langle R_i(X), C_j(Y)^T \rangle$, where $R_i(X)$ is the i th row vector of X and $C_j(Y)$ is the j th column vector of Y .]
- (d) Let M be as defined in part (b). Show that we can compute M^n in NC^2 . Deduce that $\text{NL} \subseteq \text{NC}^2$.
- (e) Strengthen the bound in part (d) to show that $\text{NL} \subseteq \text{AC}^1$.
- (f) The complexity class SAC^k is the set of languages decidable by AC^k circuits, where the AND gates have fan-in 2. Compare this to general AC^k circuits, the AND gates have unbounded fan-in. Strengthen part (e) to show that $\text{NL} \subseteq \text{SAC}^1$.

Remark 40. We have the following relations:

$$\text{NC}^0 \subsetneq \text{AC}^0 \subsetneq \text{NC}^1 \subseteq \text{L} \subseteq \text{NL} \subseteq \text{SAC}^1 \subseteq \text{AC}^1 \subseteq \text{NC}^2.$$

We have not proven that $\text{AC}^0 \neq \text{NC}^1$, nor have we shown that $\text{NC}^1 \subseteq \text{L}$.

2 Computability

Our goal in this section is to introduce the notion of Turing Machines, as well as some basic notions from Computability Theory. Computational Complexity began by trying to extend notions of Computability to solvable (decidable) problems. Rather than providing our algorithmic theorems with unlimited resources, we sought to ask which questions were decidable within given resource constraints. The most natural resource constraints to consider were time and space. While Computability Theory is very well understood, many of the analogues in Computational Complexity are not. We begin with an introduction to Turing Machines and undecidability.

These notes follow closely [Sav97][Chapter 5] and [Lev20].

2.1 Turing Machines

The Turing Machine is an abstract model of computation which has a finite set of states, as well as an infinite work tape. We assume the work tape has a left-most starting cell and is only infinite in the rightwards direction. Each cell contains a letter or is left blank. The Turing Machine has a tape head, which starts over the left-most cell. The tape head can only examine one cell at a time. Now computation steps are precisely state transitions. The Turing Machine transition function takes as input the current state and the current letter under its tape head. The transition function then does three things:

- Transitions the Turing Machine to a new state,
- Writes a new symbol to the current cell (possibly the blank symbol β , which corresponds to erasing the current letter), and
- Moves the tape head one cell, either to the left or to the right.

Intuitively, it may be helpful to think of the Turing Machine as managing an infinite, doubly-linked list data structure, where each node stores a writeable letter.

Turing Machines are used to solve decision problems. Given a language L , we ask whether it is possible to design a Turing Machine M such that for every string x , M correctly decides whether $x \in L$. To this end, the Turing Machine has explicit accept and reject states, which take effect immediately upon being reached. We will see later that not every language L can be decided in this manner. Given a Turing Machine M , we denote the set of strings that M accepts as $L(M)$.

2.1.1 Deterministic Turing Machine

We now define the standard deterministic Turing Machine as follows.

Definition 41 (Deterministic Turing Machine). A Deterministic Turing Machine is a 7-tuple

$$(Q, \Sigma, \Gamma, \delta, q_0, q_{\text{accept}}, q_{\text{reject}})$$


where Q, Σ, Γ are all finite sets and:

- Q is the set of states.
- Σ is the input alphabet, not containing the blank symbol β .
- Γ is the tape alphabet, where $\beta \in \Gamma$ and $\Sigma \subset \Gamma$. In particular, it may be helpful for the Turing Machine to have additional letters to use internally.
- $\delta : Q \times \Gamma \rightarrow Q \times \Gamma \times \{L, R\}$ is the transition function, which takes a state and tape character and returns the new state, the tape character to write to the current cell, then a direction for the tape head to move one cell to the left or right (denoted by L or R respectively).
- $q_0 \in Q$, the initial state.
- $q_{\text{accept}} \in Q$, the accept state.

- $q_{\text{reject}} \in Q$, the reject state where $q_{\text{reject}} \neq q_{\text{accept}}$.

We now consider some example of a Turing Machine.

Example 42. Let $\Sigma = \{0, 1\}$ and let $L = \{1^{2k} : k \in \mathbb{N}\} = (11)^*$. For a reader who is familiar with regular languages, observe that L is regular as L can be expressed using the regular expression $(11)^*$. A simplified computational model, known as a finite state automaton (FSM), can easily be constructed to accept L . Such a FSM diagram is provided below.



Now let's construct a Turing Machine to accept $(11)^*$. The construction of the Turing Machine, is in fact, almost identical to that of the finite state automaton. The Turing Machine will start with the input string on the far-left of the tape, with the tape head at the start of the string. The Turing Machine has $Q = \{q_0, q_1, q_{\text{reject}}, q_{\text{accept}}\}$, $\Sigma = \{0, 1\}$, and $\Gamma = \{0, 1, \beta\}$. Let the Turing Machine start at q_0 and read in the character under the tape head. If it is not a 1 or the empty string, enter q_{reject} and halt. Otherwise, if the string is empty, enter q_{accept} and halt. On the input of a 1, transition to q_1 and move the tape head one cell to the right. While in q_1 , read in the character on the tape head. If it is a 1, transition to q_0 and move the tape head one cell to the right. Otherwise, enter q_{reject} and halt. The Turing Machine always halts, and accepts the string if and only if it halts in state q_{accept} .

Observe the similarities between the Turing Machine and finite state automaton. The intuition should follow that any language accepted by a finite state automaton (ie., any regular language) can also be accepted by a Turing Machine. Formally, the Turing Machine simulates the finite state automaton by omitting the ability to write to the tape or move the tape head to the left. We now consider a second example of a Turing Machine accepting a more complicated language (namely, a context-free language).

Example 43. Let $\Sigma = \{0, 1\}$ and let $L = \{0^n 1^n : n \in \mathbb{N}\}$. For a reader familiar with automata theory, we note that L is context-free, but not regular. We provide a Turing Machine to accept this language. The Turing Machine has a tape alphabet of $\Gamma = \{0, 1, \hat{0}, \hat{1}\}$ and set of states $Q = \{q_0, q_{\text{find-1}}, q_{\text{find-0}}, q_{\text{validate}}, q_{\text{accept}}, q_{\text{reject}}\}$. Conceptually, rather than using a stack as a pushdown automaton would, the Turing Machine will use its tape head. Intuitively, the Turing Machine starts with a 0 and marks it, then moves the tape head to the right one cell at a time looking for a corresponding 1 to mark. Once it finds and marks the 1, the Turing Machine then moves the tape head to the left one cell at a time searching for the next unmarked 0 to mark. It then repeats this procedure, looking for another unmarked 1 to mark. If it finds an unpaired 0 or 1, it rejects the string. This procedure repeats until either the string is rejected, or we mark all pairs of 0's and 1's. In the latter case, the Turing Machine accepts the string.

So initially the Turing Machine starts at q_0 with the input string on the far-left of the tape, with the tape head above the first character. If the string is empty, the Turing Machine enters q_{accept} and halts. If the first character is a 1, the Turing Machine enters q_{reject} and halts. If the first character is 0, the Turing Machine replaces it with $\hat{0}$. It then moves the tape head to the right one cell and transitions to state $q_{\text{find-1}}$.

At state $q_{\text{find-1}}$, the Turing Machine moves the tape head to the right and stays at $q_{\text{find-1}}$ for each 0 or $\hat{0}$ character it reads in and writes back the character it parsed. If at $q_{\text{find-1}}$ and the Turing Machine reads 1, then it writes $\hat{1}$ to the tape, moves the tape head to the left, and transitions to $q_{\text{find-0}}$. If no 1 is found, the Turing Machine enters q_{reject} and halts.

At state $q_{\text{find-0}}$, the Turing Machine moves the tape head to the left and stays at $q_{\text{find-0}}$ until it reads in 0. If the Turing Machine reads in 0 at state $q_{\text{find-0}}$, it replaces the 0 with $\hat{0}$. It then moves the tape head to the

right one cell and transitions to state $q_{\text{find-1}}$. If no 0 is found once we have reached the far-left cell, the Turing Machine transitions to state q_{validate} .

At state q_{validate} , the Turing Machine transitions to the right one cell at a time while staying at q_{validate} . If it encounters any 1, it enters q_{reject} . Otherwise, the Turing Machine enters q_{accept} once reading in β .

Remark 44. Now that we provided formal specifications for a couple Turing Machines, we provide a more abstract representation from here on out. We are more interested in studying the power of Turing Machines rather than the individual state transitions, so high level procedures suffice for our purposes. This high level procedure from the above example provides sufficient detail to simulate the Turing Machine. So for our purposes, this level of detail is sufficient:

“Intuitively, the Turing Machine starts with a 0 and marks it, then moves the tape head to the right one cell at a time looking for a corresponding 1 to mark. Once it finds and marks the 1, the Turing Machine then moves the tape head to the left one cell at a time searching for the next unmarked 0 to mark. It then repeats this procedure, looking for another unmarked 1 to mark. If it finds an unpaired 0 or 1, it rejects the string. This procedure repeats until either the string is rejected, or we mark all pairs of 0’s and 1’s. In the latter case, the Turing Machine accepts the string.”

Definition 45 (Recursively Enumerable Language). A language L is said to be *recursively enumerable* if there exists a deterministic Turing Machine M such that $L(M) = L$. Note that if $\omega \notin L$, the machine M need not halt on ω .

Definition 46 (Decidable Language). A language L is said to be *decidable* if L is there exists some Turing Machine M such that $L(M) = L$ and M halts on all inputs. We say that M *decides* L .

Remark 47. Every decidable language is clearly recursively enumerable. The converse is not true, and this will be shown later with the undecidability of the Halting problem.

2.1.2 Multitape Turing Machine

The Turing Machine is quite a robust model, in the sense that the standard deterministic model accepts and decides precisely the same languages as the multitape and non-deterministic variants. It should be noted that one model may actually be more efficient than another. In regards to language acceptance and computability, we ignore issues of efficiency and complexity. However, the same techniques we use to show that these models are equally powerful can be leveraged to show that two models of computation are equivalent both with regards to power and some measure of efficiency. That is, to show that the two models solve the same set of problems using a comparable amount of resources (e.g., polynomial time). This is particularly important in complexity theory, but we also leverage these techniques when showing Turing Machines equivalent to other models such as (but not limited to) the RAM model and the λ -calculus.

We begin by introducing the Multitape Turing Machine.

Definition 48 (Multitape Turing Machine). A *k-tape Turing Machine* is an extension of the standard deterministic Turing Machine in which there are k tapes with infinite memory and a fixed beginning. The input initially appears on the first tape, starting at the far-left cell. The transition function is the addition difference, allowing the k -tape Turing Machine to simultaneously read from and write to each of the k -tapes, as well as move some or all of the tape cells. Formally, the transition function is given below:

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

The expression:

$$\delta(q_i, a_1, \dots, a_k) = (q_j, b_1, \dots, b_k, L, R, S, \dots, R)$$

indicates that the TM is on state q_i , reading a_m from tape m for each $m \in [k]$. Then for each $m \in [k]$, the TM writes b_m to the cell in tape m highlighted by its tape head. The m th component in $\{L, R, S\}^k$ indicates that the m th tape head should move left, right, or remain stationary respectively.

Our first goal is to show that the standard deterministic Turing Machine is equally as powerful as the k -tape Turing Machine, for any $k \in \mathbb{N}$. We need to show that the languages accepted (decided) by deterministic Turing Machines are exactly those languages accepted (decided) by the multitape variant. The initial approach of

a set containment argument is correct. The details are not as intuitively obvious. Formally, we show that for any multitape Turing Machine, there exists an deterministic Turing Machine; and for any deterministic Turing Machine, there exists an equivalent multitape Turing Machine. In other words, we show how one model simulates the other and vice-versa. This implies that the languages accepted (decided) by one model are precisely the languages accepted (decided) by the other model.

Theorem 49. A language is recursively enumerable (decidable) if and only if some multitape Turing Machine accepts (decides) it.

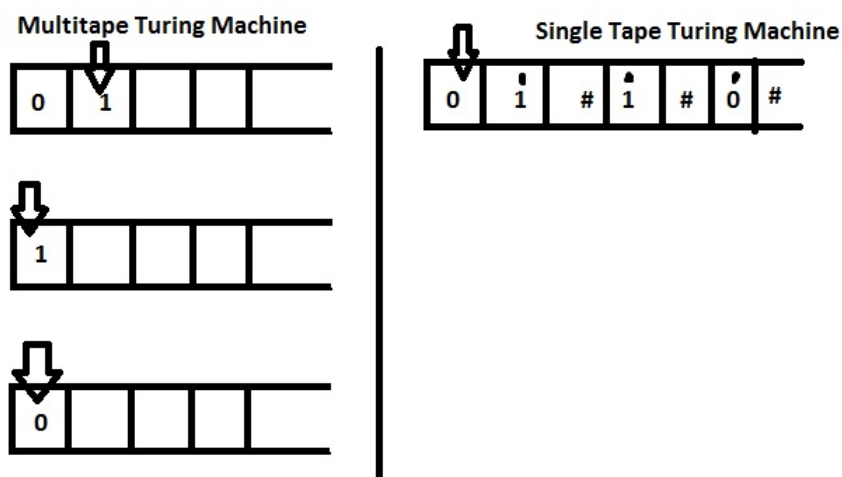
Proof. We begin by showing that the multitape Turing Machine model is at least as power as the standard deterministic Turing Machine model. Clearly, a standard deterministic Turing Machine is a 1-tape Turing Machine. So every language accepted (decided) by a standard deterministic Turing Machine is also accepted (decided) by some multitape Turing Machine.

Conversely, let M be a multitape Turing Machine with k tapes. We construct a standard deterministic Turing Machine M' to simulate M , which shows that $L(M') = L(M)$. As M has k -tapes, it is necessary to represent the strings on each of the k tapes on a single tape. It is also necessary to represent the placement of each of the k tape heads of M on the one tape of M' . This is done by using a special marker. For each symbol $c \in \Gamma(M)$, we include c and \hat{c} in Γ' , where \hat{c} indicates a tape head on M is on the character c . We then have a special delimiter symbol $\#$, which separates the strings on each of the k tapes. So $|\Gamma(M')| = 2|\Gamma(M)| + 1$. M' simulates M in the following manner.

- M' scans the tape from the first delimiter to the last delimiter to determine which symbols are marked as under the tape heads on M .
- M' then evaluates the transition function of M , then makes a second pass along the tape to update the symbols on the tape.
- If at any point, the tape head of M' falls on a delimiter symbol $\#$, M would have reached the end of that specific tape. So M' shifts the string, cell by cell, starting at the current delimiter inclusive. A blank symbol is then overwritten on the delimiter.

Thus, M' simulates M . So any language accepted (decided) by a multitape Turing Machine is accepted (decided) by a single tape Turing Machine. □

Below is an illustration of a multitape Turing Machine and an equivalent single tape Turing Machine.



Remark 50. If the k -tape Turing Machine takes T steps, then each tape uses at most $T + 1$ cells. So the equivalent one-tape deterministic Turing Machine constructed in the proof of Theorem 4.1 takes $(k \cdot (T + 1))^2 = \mathcal{O}(T^2)$ steps.

2.1.3 Non-deterministic Turing Machines

We now introduce the non-deterministic Turing Machine. The non-deterministic Turing Machine has a single, infinite tape with an end at the far-left. Its sole difference with the deterministic Turing Machine is the transition function. For a given state, letter pair $(q, a) \in Q \times \Gamma$, the transition function evaluated at this pair $\delta(q, a)$ specifies a unique next state. For a non-deterministic Turing Machine, there may be multiple permissible states to visit next. We make this precise in the definition by defining the transition function to return a subset of elements from $Q \times \Gamma \times \{L, R\}$. For a computation, the non-deterministic Turing Machine makes a selection of which of the possible transitions to consider from the permissible choices. With this in mind, we formalize the non-deterministic Turing Machine.

Definition 51 (Non-deterministic Turing Machine). A Non-Deterministic Turing Machine is a 7-tuple

$$(Q, \Sigma, \Gamma, \delta, q_0, q_{\text{accept}}, q_{\text{reject}})$$

where Q, Σ, Γ are all finite sets and:

- Q is the set of states.
- Σ is the input alphabet, not containing the blank symbol β .
- Γ is the tape alphabet, where $\beta \in \Gamma$ and $\Sigma \subset \Gamma$. In particular, it may be helpful for the Turing Machine to have additional letters to use internally.
- $\delta : Q \times \Gamma \rightarrow 2^{Q \times \Gamma \times \{L, R\}}$ is the transition function, which takes a state and tape character and returns the new state, the tape character to write to the current cell, then a direction for the tape head to move one cell to the left or right (denoted by L or R respectively).
- $q_0 \in Q$, the initial state.
- $q_{\text{accept}} \in Q$, the accept state.
- $q_{\text{reject}} \in Q$, the reject state where $q_{\text{reject}} \neq q_{\text{accept}}$.

We now show that the deterministic and non-deterministic variants are equally powerful. The proof for this is by simulation. Before introducing the proof, let's conceptualize this. Earlier in this section, a graph theory intuition was introduced for understanding the definition of what it means for a string to be accepted by a non-deterministic Turing Machine. That definition of string acceptance dealt with the existence of a choice string such that the non-deterministic Turing Machine would reach the accept state q_{accept} from the starting state q_0 . The graph theory analog was that there existed a path for the input string from q_0 to q_{accept} .

So the way a deterministic Turing Machine simulates a non-deterministic Turing Machine is through, essentially, a breadth-first search. More formally, what actually happens is that a deterministic multitape Turing Machine is used to simulate a non-deterministic Turing Machine. It does this by generating choice strings in lexicographic order and simulating the non-deterministic Turing Machine on each choice string until the string is accepted or all the possibilities are exhausted.

Given the non-deterministic Turing Machine has a finite number of transitions, there are a finite number of choice input strings to generate. Thus, a multitape deterministic Turing Machine will always be able to determine if an input string is accepted by the non-deterministic Turing Machine. It was already proven that a multitape Turing Machine can be simulated by a standard deterministic Turing Machine, so it follows that any language accepted by a non-deterministic Turing Machine can also be accepted by a deterministic Turing Machine.

Theorem 52. A language is recursively enumerable (decidable) if and only if it is accepted (decided) by some non-deterministic Turing Machine.

Proof. A deterministic Turing Machine is clearly non-deterministic. So it suffices to show that every non-deterministic Turing Machine has an equivalent deterministic Turing Machine. From Theorem 49, it suffices to construct a multitape Turing Machine equivalent for every non-deterministic Turing Machine. The proof is

by simulation. Let M be a non-deterministic Turing Machine.

We construct a three-tape Turing Machine M' to simulate all possibilities. The first tape contains the input string and is used as a read-only tape. The second tape is used to simulate M , and the third tape is the enumeration tape in which we enumerate the branches of the non-deterministic Turing Machine. Let:

$$b = \max_{q \in Q, a \in \Gamma} |\delta_M(q, a)|.$$

The tape alphabet of M' is $\Gamma(M) \cup [b]$. On the third tape of M' , we enumerate strings over $[b]^n$ in lexical order, where n is the length of the input string. At state i in the computation, we utilize the transition indexed by the number on the i th cell on the third tape.

Formally, M' works as follows:

1. M' is started with the input ω on the first tape.
2. We then copy the input string to the second tape and generate 0^ω .
3. Simulate M on ω using the choice string on ω . If at any point, the transition specified by the third tape is undefined (which may occur if too few choices are available), we terminate the simulation of M and generate the next string in lexical order on the third tape. We then repeat step (3).
4. M' accepts (rejects) ω if and only some simulation of M on ω accepts (rejects) ω .

By construction, $L(M') = L(M)$, yielding the desired result. □

2.1.4 Exercises

(Recommended) Problem 21. Let $L = \{a^n b^{2n} : n \in \mathbb{Z}^+\}$. Provide a high level description for a Turing Machine that accepts L . Your description may suppress individual state transitions, but it should be detailed enough to allow you (or one of your classmates) to construct the exact transition function. In particular, your description should specify the movement of the tape head and when the Turing Machine writes to the tape.

2.2 Undecidability

In this section, we examine the limits of computation as a means to solve problems. This is important for several reasons. First, problems that cannot be solved need to be simplified to a formulation that is more amenable to computational approaches. Second, the techniques used in proving languages to be undecidable, including reductions and diagonalization, appear repeatedly in complexity theory. Lastly, undecidability is an interesting topic in its own right.

The canonical result in computability theory is the undecidability of the halting problem. Intuitively, no algorithm exists to decide if a Turing Machine halts on an arbitrary input string. While the result seems abstract and unimportant, the results are actually far reaching. Software engineers seek better ways to determine the correctness of their programs. The undecidability of the halting problem provides an impossibility result for software engineers; no such techniques exist to validate arbitrary computer programs. We formalize this with the following language:

$$A_{\text{TM}} = \{\langle M, w \rangle : M \text{ is a Turing Machine that accepts } w\}.$$

It turns out that A_{TM} is undecidable. We actually start by showing that the following language is undecidable:

$$L_{\text{diag}} = \{\omega_i : \omega_i \text{ is the } i\text{th string in } \Sigma^*, \text{ which is accepted by the } i\text{th Turing Machine } M_i\}.$$

L_{diag} is designed to leverage a diagonalization argument. We note that Turing Machines are representable as finite strings (just like computer programs), and that the set of finite length strings over an alphabet is countable. So we can enumerate Turing Machines using \mathbb{N} . Similarly, we also enumerate input strings from Σ^* using \mathbb{N} . Before proving L_{diag} undecidable, we need to show that decidable languages are precisely those that are recursively enumerable and whose complements are recursively enumerable. We unpack this idea before proving the theorem.

Definition 53. Let REC be the set of decidable languages, and let RE be the set of recursively enumerable languages. Denote coRE to be the set of languages L , where $\bar{L} \in \text{RE}$.

Theorem 54. A language L is decidable if and only if L and \bar{L} are recursively enumerable. That is, $\text{REC} = \text{RE} \cap \text{coRE}$.

Proof. Suppose first L is decidable, and let M be a decider for L . As M decides L , M also accepts L . So L is recursively enumerable. Now define \bar{M} to be a Turing Machine that, on input ω , simulates M on ω . \bar{M} accepts (rejects) ω if and only if M rejects (accepts) ω . As M is a decider, \bar{M} decides \bar{L} . So \bar{L} is also recursively enumerable.

Conversely, suppose L and \bar{L} are recursively enumerable. Let B and \bar{B} be Turing Machines that accept L and \bar{L} respectively. We construct a Turing Machine K to decide L . K works as follows. On input ω , K simulates B and \bar{B} in parallel on ω . As L and \bar{L} are recursively enumerable, at least one of B or \bar{B} will halt and accept ω . If B accepts ω , then so does K . Otherwise, \bar{B} accepts ω and K rejects ω . So K decides L . \square

Remark 55. While $\text{RE} \cap \text{coRE}$ is very well understood, the analogues in Computational Complexity remain long-standing open problems. As an example, it is unknown whether $\text{P} = \text{NP} \cap \text{coNP}$.

In order to show that L_{diag} is undecidable, we have by Theorem 54 that it suffices to show $\overline{L_{\text{diag}}}$ is not recursively enumerable. This is the meat of the proof for the undecidability of the halting problem. It turns out that L_{diag} is recursively enumerable, which is easy to see.

Theorem 56. L_{diag} is recursively enumerable.

Proof. We construct an acceptor D for L_{diag} which works as follows. On input ω_i , D simulates M_i on ω_i and accepts ω_i if and only if M_i accepts ω_i . So $L(D) = L_{\text{diag}}$, and L_{diag} is recursively enumerable. \square

We now show that $\overline{L_{\text{diag}}}$ is not recursively enumerable.

Theorem 57. $\overline{L_{\text{diag}}}$ is not recursively enumerable.

Proof. Suppose to the contrary that $\overline{L_{\text{diag}}}$ is recursively enumerable. Let $k \in \mathbb{N}$ such that the Turing Machine M_k accepts $\overline{L_{\text{diag}}}$. Suppose $\omega_k \in \overline{L_{\text{diag}}}$. Then M_k accepts ω_k , as $L(M_k) = \overline{L_{\text{diag}}}$. However, $\omega_k \in \overline{L_{\text{diag}}}$ implies that M_k does not accept ω_k , a contradiction.

Suppose instead $\omega_k \notin \overline{L_{\text{diag}}}$. Then $\omega_k \notin L(M_k) = \overline{L_{\text{diag}}}$. Since M_k does not accept ω_k , it follows by definition of $\overline{L_{\text{diag}}}$ that $\omega_k \in \overline{L_{\text{diag}}}$, a contradiction. So $\omega_k \in \overline{L_{\text{diag}}}$ if and only if $\omega_k \notin \overline{L_{\text{diag}}}$. So $\overline{L_{\text{diag}}}$ is not recursively enumerable. \square

Corollary 58. L_{diag} is undecidable.

Proof. This follows immediately from Theorem 54, as L_{diag} is recursively enumerable, while $\overline{L_{\text{diag}}}$ is not. \square

2.3 Reducibility

The goal of a reduction is to transform one problem A into another problem B . If we know how to solve this second problem B , then this yields a solution for A . Essentially, we transform A into B , solve it in B , then apply this solution in A . Reductions thus allow us to order problems based on how hard they are. In particular, if we know that A is undecidable, a reduction immediately implies that B is undecidable. Otherwise, a Turing Machine to decide B could be used to decide A . Reductions are also a standard tool in complexity theory, where we transform problems with some bound on resources (such as time or space bounds). In computability theory, reductions need only be computable. We formalize the notion of a reduction with the following two definitions.

Definition 59 (Computable Function). A function $f : \Sigma^* \rightarrow \Sigma^*$ is a *computable function* if there exists some Turing Machine M such that on input ω , M halts with just $f(\omega)$ written on the tape.

Definition 60 (Many-to-One Reduction). Let A, B be languages. A *many-to-one reduction* from A to B is a computable function $f : \Sigma^* \rightarrow \Sigma^*$ such that $\omega \in A$ if and only if $f(\omega) \in B$. We say that A is reducible to B , denoted $A \leq_m B$, if there exists a many-to-one reduction from A to B .

We deal with reductions in a similar high-level manner as Turing Machines, providing sufficient detail to indicate how the original problem instances are transformed into instances of the target problem. In order for reductions to be useful in computability theory, we need an initial undecidable problem. This is the language L_{diag} from the previous section. With the idea of a reduction in mind, we proceed to show that A_{TM} is undecidable.

Theorem 61. A_{TM} is undecidable.

Proof. It suffices to show $L_{\text{diag}} \leq_m A_{\text{TM}}$. The function $f : \Sigma^* \rightarrow \Sigma^*$ maps $\omega_i \in L_{\text{diag}}$ to $\langle M_i, \omega_i \rangle \in A_{\text{TM}}$. Any string not in L_{diag} is mapped to ϵ under f . As Turing Machines are enumerable, a Turing Machine can clearly write $\langle M_i, \omega_i \rangle$ to the tape when started with ω_i . So f is computable. Furthermore, observe that $\omega_i \in L_{\text{diag}}$ if and only if $\langle M_i, \omega_i \rangle \in A_{\text{TM}}$. So f is a reduction from L_{diag} to A_{TM} and we conclude that A_{TM} is undecidable. \square

With A_{TM} in tow, we prove the undecidability of the halting problem, which is given by:

$$H_{\text{TM}} = \{ \langle M, w \rangle : M \text{ is a Turing Machine that halts on the string } w \}.$$

Theorem 62. H_{TM} is undecidable.

Proof. It suffices to show that $A_{\text{TM}} \leq_m H_{\text{TM}}$. Each element of A_{TM} is clearly an element of H_{TM} . So we map each element of A_{TM} to itself in H_{TM} , and all other strings to ϵ . This map is clearly a reduction, so H_{TM} is undecidable. \square

The reductions to show A_{TM} and H_{TM} undecidable have been rather trivial. We will examine some additional undecidable problems. In particular, the reduction will be from A_{TM} . The idea moving forward is to pick a desirable solution and return it if and only if a Turing Machine M halts on a string ω . A decider for the target problem would thus give us a decider for A_{TM} , which is undecidable. We illustrate the concept below.

Theorem 63. Let $E_{\text{TM}} = \{ \langle M \rangle : M \text{ is a Turing Machine s.t. } L(M) = \emptyset \}$. E_{TM} is undecidable.

Proof. It suffices to show that $A_{\text{TM}} \leq_m E_{\text{TM}}$. For each instance of $\langle M, \omega \rangle \in A_{\text{TM}}$, we construct an instance of E_{TM} M' as follows. On input $x \neq \omega$, M' rejects x . Otherwise, M' simulates M on ω . If M accepts (rejects) ω , then M' rejects (accepts) ω . So $\langle M, \omega \rangle \in A_{\text{TM}}$ implies that $M' \in E_{\text{TM}}$. So E_{TM} is undecidable. \square

We use the same idea to show that it is undecidable if a Turing Machine accepts the empty string. Observe above that our desirable solution for E_{TM} was \emptyset . Then M' accepted the desired solution if and only if the instance Turing Machine M accepted ω . We *conditioned* acceptance of the target instance based on the original problem. In this next problem, the target solution is ϵ , the empty string.

Theorem 64. Let $L_{\text{ES}} = \{\langle M \rangle : M \text{ is a Turing Machine that accepts } \epsilon\}$. L_{ES} is undecidable.

Proof. We show that $A_{\text{TM}} \leq_m E_{\text{TM}}$. Let $\langle M, \omega \rangle \in A_{\text{TM}}$. We construct an instance of L_{ES} , M' , as follows. On input $x \neq \epsilon$, M' rejects x . Otherwise, M' simulates M on ω . M' accepts ϵ if and only if M accepts ω . So $\langle M, \omega \rangle \in A_{\text{TM}}$ if and only if $M' \in L_{\text{ES}}$. This function is clearly computable, so L_{ES} is undecidable. \square

Recall that any regular language is decidable. We may similarly ask if a given language is regular. It turns out that this new problem is undecidable.

Theorem 65. Let $L_{\text{Reg}} = \{L : L \text{ is regular}\}$. L_{Reg} is undecidable.

Proof. We reduce A_{TM} to L_{Reg} . Let $\langle M, \omega \rangle \in A_{\text{TM}}$. We construct a Turing Machine M' such that $L(M')$ is regular if and only if M accepts ω . M' works as follows. On input x , M' accepts x if it is of the form $0^n 1^n$ for some $n \in \mathbb{N}$. Otherwise, M' simulates M on ω , and accepts x if and only if M accepts ω . So $L(M') = \Sigma^*$ if and only if M accepts ω , and $L(M') = \{0^n 1^n : n \in \mathbb{N}\}$ otherwise which is not regular. Thus, $L(M') \in L_{\text{Reg}}$ if and only if $\langle M, \omega \rangle \in A_{\text{TM}}$. So L_{Reg} is undecidable. \square

The common theme in each of these undecidability results is that not every language satisfies the given property. This leads us to one of the major results in computability theory: Rice's Theorem. Intuitively, Rice's Theorem states that any non-trivial property is undecidable. A property is said to be trivial if it applies to either every language or no language. We formalize it as follows.

Theorem 66 (Rice). Let C be a non-empty, proper subset of RE. Then C is undecidable.

Proof. We reduce A_{TM} to C . Without loss of generality, suppose $\emptyset \in C$. As C is a proper subset of RE, \overline{C} is non-empty. Let $L \in \overline{C}$. Let $\langle M, \omega \rangle \in A_{\text{TM}}$. We construct a Turing Machine M' as follows. On input x , M' rejects x if $x \notin L$. Otherwise, M' simulates M on ω . M' rejects x if and only if M accepts ω . So $L(M') = \emptyset \in C$ if and only if $\langle M, \omega \rangle \in A_{\text{TM}}$. Otherwise, $L(M') = L$. Thus, C is undecidable. \square

Remark 67. Observe that the proof of Rice's Theorem is a template for the previous undecidability proofs in this section. Rice's Theorem generalizes all of our undecidability results and provides an easy test to determine if a language is undecidable. In short, to show a property undecidable, it suffices to exhibit a language satisfying said property and a language that does not satisfy said property.

2.3.1 Exercises

(Recommended) Problem 22. Consider the language:

$$L = \{\langle M, \omega \rangle : M \text{ halts in at most 2020 steps when run on } \omega\}.$$

Do the following.

- (a) Show that L is recursively enumerable.
- (b) Give a many-to-one reduction from Halt_{TM} to L . Conclude that L is undecidable.

(Recommended) Problem 23. *Fermat's Last Theorem* states that there are no positive integer solutions to the equation $x^n + y^n = z^n$ for any integer $n > 2$. This theorem was conjectured by Fermat in 1637 and finally proven in 1994 by Andrew Wiles. Consider the following Turing Machine M , which searches for counterexamples to Fermat's Last Theorem:

" M ignores any input (so it can be assumed to run when started on a blank tape), and proceeds to enumerate 4-tuples of the form (x, y, z, n) where x, y, z, n are positive integers and $n > 2$. After enumerating each 4-tuple,

M computes x^n, y^n, z^n , then checks that $x^n + y^n = z^n$ and $n > 2$. M halts if $x^n + y^n = z^n$ and $n > 2$. Otherwise, M proceeds to the next 4-tuple.”

Explain how an algorithm to decide the Halting Problem could be used to decide if Fermat’s Last Theorem were true. [**Note:** You do not need to know anything about Fermat’s Last Theorem to answer this question.]

(Recommended) Problem 24. Let \mathcal{C} be the set of recursively enumerable languages that are co-finite. That is, $L \in \mathcal{C}$ precisely if \bar{L} is finite. Using Rice’s Theorem, show that \mathcal{C} is undecidable.

2.4 Oracle Turing Machines

In this section, we introduce Oracle Turing Machines. The key idea is as follows. Suppose we are trying to solve a computational problem L . Given access to an algorithm A to solve a different (possibly harder) problem L' , can we use A to solve L ? If so, how many queries to A are required? Note that we make no assumptions that L' is decidable, or even recursively enumerable. In this sense, we can make precise bottlenecks or barriers in solving computational problems. In particular, measuring the number of queries to a given oracle, known as the *query complexity*, is a very active area in Computational Complexity (particularly in Quantum Complexity). It is quite possible that one has come across examples of this already in an Algorithms or Theory of Computation course, in reducing a search problem to the corresponding decision problem. For example, one may have seen already how to use a SAT oracle to find explicit satisfying instances for Boolean formulas.

We begin with the definition of an Oracle.

Definition 68. Let Σ be an alphabet. An *oracle* is a set $\mathcal{O} \subset \Sigma^*$.

We next introduce the Oracle Turing Machine model. Informally, an Oracle Turing Machine M is a Turing Machine with the ability to query an oracle. We think of M as separate from the underlying oracle. That is, the transition function is associated to M and is independent of the oracle that is used. While the transition function is fixed, whether M accepts the input string ω depends significantly on the underlying oracle. Suppose we have the Oracles \mathcal{O}_1 and \mathcal{O}_2 . On input ω , M poses an Oracle query to ask whether the string x is in the Oracle. If $x \in \mathcal{O}_1$ and $x \notin \mathcal{O}_2$, then M may behave differently based on which oracle is used. In particular, equipping M with \mathcal{O}_1 (which we denote $M^{\mathcal{O}_1}$) might result in $M^{\mathcal{O}_1}(\omega) = 1$, while $M^{\mathcal{O}_2}(\omega) = 0$.

Definition 69. An *Oracle Turing Machine* M is a Turing Machine, which is parameterized by an Oracle \mathcal{O} . M with three additional states: q_{query} , q_{yes} , and q_{no} , as well as an additional tape to which M has both read and write access. The Turing Machine may query the Oracle by entering the state q_{query} and writing a string x to the additional tape. If x is in the underlying Oracle, then M transitions to q_{yes} . Otherwise, M transitions to q_{no} . Querying the Oracle takes a single computation step.

With the notion of an Oracle Turing Machine in tow, we can now begin talking about complexity classes relative to Oracles. We refer to such classes as *relativized*. Note as well that some care needs to be taken, in that we can only define relativized complexity classes that have characterizations in terms of Turing Machines. We have not discussed how to define relativized complexity classes where the underlying model of computation is a circuit, for instance.

Definition 70. Let \mathcal{C} be a complexity class characterized by Turing Machines, and let \mathcal{O} be an Oracle. A language $L \in \mathcal{C}^{\mathcal{O}}$ if there exists a \mathcal{C} -Turing Machine M with oracle \mathcal{O} , denoted $M^{\mathcal{O}}$, such that $M^{\mathcal{O}}$ accepts L .

Example 71. As an example, the language $L \in \text{P}^{\text{SAT}}$ if there exists a polynomial-time deterministic, Oracle Turing Machines M^{SAT} that decides L .

We may want to talk about relativizing a complexity class \mathcal{A} with respect to a second complexity class \mathcal{B} , as opposed to just an individual problem. This brings us to the following definition.

Definition 72. Let \mathcal{A} and \mathcal{B} be complexity classes defined in terms of Turing Machines. Define:

$$\mathcal{A}^{\mathcal{B}} := \bigcup_{L \in \mathcal{B}} \mathcal{A}^L.$$

Remark 73. We note that if \mathcal{B} has a complete problem L under \mathcal{A} -computable reductions, then $\mathcal{A}^{\mathcal{B}} = \mathcal{A}^L$. We organize this statement with the next theorem, the proof of which is left as an exercise.

Theorem 74. Let \mathcal{A} and \mathcal{B} be complexity classes defined in terms of Turing Machines. Suppose that \mathcal{B} has a complete problem L , under \mathcal{A} -computable many-to-one reductions. Then $\mathcal{A}^L = \mathcal{A}^{\mathcal{B}}$.

Example 75. Take $\mathcal{A} = \text{P}$, $\mathcal{B} = \text{NP}$, and $L = \text{SAT}$. We note that NP-complete problems are P-computable. So by Theorem 74 $\text{P}^{\text{NP}} = \text{P}^{\text{SAT}}$.

Example 76. Take $\mathcal{A} = \text{L}$, $\mathcal{B} = \text{NP}$, and $L = \text{SAT}$. We note that NP-complete problems are P-computable. While L-computable functions are computable in P, the converse is not known to be true. So Theorem 74 is not known to apply.

Example 77. Take $\mathcal{A} = \text{P}$ and $\mathcal{B} = \text{BPP}$.⁴ It remains open as to whether BPP has complete problems. So under current knowledge, Theorem 74 is not applicable in this setting.

Now that we have some handle on the notion of an Oracle, we conclude with the notion of a Turing reduction. Informally, we say that a language L_1 is *Turing reducible* to L_2 if given an algorithm to solve L_2 , we can solve L_1 . The idea of providing an algorithm is formalized with the notion of Oracle Turing Machines.

Definition 78 (Turing Reduction). Let L_1 and L_2 be languages. We say that L_1 is *Turing reducible* to L_2 , denoted $L_1 \leq_T L_2$, if there exists an Oracle Turing Machine M^{L_2} that decides L_1 .

Remark 79. Recall that the notion of a Turing reduction is how we intuitively explain why the many-to-one Karp reductions indeed establish hardness. Suppose that $L_1 \leq_m L_2$, and let $\varphi : \Sigma^* \rightarrow \Sigma^*$ be a many-to-one Karp reduction. Suppose that we can solve L_2 . Let $\omega \in \Sigma^*$. In order to decide if $\omega \in L_1$, we compute $\varphi(\omega)$ and use the algorithm for L_2 to test whether $\varphi(\omega) \in L_2$. As φ is a reduction, we conclude that $\omega \in L_1$ if and only if the algorithm for L_2 returns that $\varphi(\omega) \in L_2$. So in fact, the fact that L_1 is many-to-one Karp reducible to L_2 implies that L_1 is Turing reducible to L_2 .

2.4.1 Exercises

(Recommended) Problem 25. Do the following.

- (a) Show that $\text{NP} \subseteq \text{P}^{\text{NP}}$.
- (b) Show that $\overline{\text{SAT}} \in \text{P}^{\text{NP}}$. [**Hint:** Use the fact that $\text{P}^{\text{NP}} = \text{P}^{\text{SAT}}$.]
- (c) Deduce that $\text{coNP} \subseteq \text{P}^{\text{NP}}$.

(Recommended) Problem 26. Let \mathcal{A} and \mathcal{B} be complexity classes defined in terms of Turing Machines. Suppose that \mathcal{B} has a complete problem L , under \mathcal{A} -computable reductions. Prove that $\mathcal{A}^L = \mathcal{A}^{\mathcal{B}}$.

Definition 80. Let \mathcal{A} and \mathcal{B} be complexity classes defined in terms of Turing Machines. We say that \mathcal{A} is *low* for \mathcal{B} if $\mathcal{B} = \mathcal{B}^{\mathcal{A}}$. That is, access to an \mathcal{A} -oracle does not increase the power of \mathcal{B} -Turing Machines.

(Recommended) Problem 27. Recall that REC is the set of decidable languages, and RE is the set of recursively enumerable languages. Do the following.

- (a) Show that REC is low for itself. That is, show that $\text{REC} = \text{REC}^{\text{REC}}$.
- (b) Show that REC is low for RE. That is, show that $\text{RE} = \text{RE}^{\text{REC}}$.

(Recommended) Problem 28. Let L_1 and L_2 be languages. Suppose that L_1 is many-to-one (Karp) reducible to L_2 ; that is, $L_1 \leq_m L_2$. Show that L_1 is Turing reducible to L_2 ; that is, $L_1 \leq_T L_2$.

⁴A language L is in BPP if there exists a randomized Turing Machine M such that for all $x \in L$, $M(x) = 1$ with probability at least $2/3$; and for all $x \notin L$, $M(x) = 1$ with probability at most $1/3$. Here, BPP stands for *bounded-error probabilistic polynomial time*.

2.5 Arithmetic Hierarchy

In this section, we introduce the Arithmetic Hierarchy. A key idea behind Oracle Turing Machines is as follows: given the ability to solve a specific problem, what additional problems can we solve? It is natural to ask as to the problems we could solve, under the assumption that the Halting problem was solvable. Similarly, are there problems that cannot be solved, even in the presence of an Oracle to the Halting problem? These questions motivate the Arithmetic Hierarchy.

We begin by denoting $\Sigma_1^0 = \text{RE}$ and $\Pi_1^0 = \text{co}\Sigma_1^0 = \text{coRE}$. Denote $\Delta_1^0 := \Sigma_1^0 \cap \Pi_1^0$. So $\Delta_1^0 = \text{REC}$. We now define the Arithmetic Hierarchy.

Definition 81. For $i > 1$, denote $\Sigma_i^0 = \text{RE}^{\Sigma_{i-1}^0}$ and $\Pi_i^0 = \text{co}\Sigma_i^0$. Denote $\Delta_i^0 = \Sigma_i^0 \cap \Pi_i^0$.

Remark 82. We note that Π_i^0 also has a characterization in terms of Oracles. Namely, $\Pi_i^0 = \text{coRE}^{\Pi_{i-1}^0}$. We will not formally prove this equivalence.

Definition 83 (Arithmetic Hierarchy). The *Arithmetic Hierarchy*, denoted AH is:

$$\text{AH} = \bigcup_{n \in \mathbb{Z}^+} \Sigma_n^0.$$

We now ask whether all the Σ_i^0, Π_i^0 , and Δ_i^0 classes are distinct. It turns out that $\Sigma_i^0 \neq \Pi_i^0$ (see Problem 29). We next introduce the following theorem, which establishes clear containment relations amongst the Σ_i^0 and Π_i^0 classes. Hence, we have a better handle on the hierarchy.

Theorem 84. For each $i \geq 1$, we have that $\Sigma_i^0 \subseteq \Delta_{i+1}^0$ and $\Pi_i^0 \subseteq \Delta_{i+1}^0$.

We leave the proof as an exercise. It follows from Theorem 84 and Problem 30 that for all $i \geq 1$, $\Sigma_i^0 \subsetneq \Sigma_{i+1}^0$ and $\Pi_i^0 \subsetneq \Pi_{i+1}^0$. So AH does not collapse to any level i .

2.5.1 Exercises

(Recommended) Problem 29. Let \mathcal{O} be an arbitrary oracle. Do the following.

- Show that $\text{REC}^{\mathcal{O}}$ is closed under complements.
- Show that $\text{REC}^{\mathcal{O}} = \text{RE}^{\mathcal{O}} \cap \text{coRE}^{\mathcal{O}}$.
- Exhibit a language $L \in \text{RE}^{\mathcal{O}}$ such that $\bar{L} \notin \text{RE}^{\mathcal{O}}$. Note that as $L \in \text{RE}^{\mathcal{O}}$, $\bar{L} \in \text{coRE}^{\mathcal{O}}$. [**Hint:** Mimic the proof techniques in the Undecidability section of the Computability notes.]
- In light of part (c), conclude that there exist problems that do not belong to $\text{REC}^{\mathcal{O}}$. That is, for every oracle \mathcal{O} , there exist undecidable problems relative to \mathcal{O} .

Remark 85. Observe that the proofs in the non-relativized setting (i.e., when dealing with Turing Machines that do not have oracles) went through immediately and were not sensitive to the presence of oracles. This is a phenomenon in Computability Theory that does not happen in Computational Complexity. The P vs. NP problem, for instance, is sensitive to the presence of oracles. As a result, Computability techniques are insufficient to resolve the P vs. NP problem. This barrier, known as the *Relativization Barrier*, is a celebrated result from 1975 due to Baker, Gill, and Solovay.

(Recommended) Problem 30. Fix $i \geq 1$. Show the following.

- $\Sigma_i^0 \subseteq \Delta_{i+1}^0$.
- $\Pi_i^0 \subseteq \Delta_{i+1}^0$.

3 Structural Complexity

Computational Complexity takes a different perspective than the study of Algorithms. We seek to classify problems into classes according to whether they can be solved within specified resource bounds. Structural Complexity Theory seeks to examine the relationships between these complexity classes. We assume familiarity with the classes P and NP, as well as NP-completeness.

3.1 Ladner's Theorem

One key approach to resolving the P vs. NP problem is to find a problem $L \in \text{NP}$ such that $L \notin \text{P}$ and L is not NP-complete. We refer to such languages as NP-intermediate. Showing that there exists an NP-intermediate language would show that $\text{P} \neq \text{NP}$. We now consider the converse. There are three possibilities to resolve the P vs. NP problem. First, suppose that $\text{P} = \text{NP}$. Then there are no NP-intermediate languages. Suppose instead that $\text{P} \neq \text{NP}$. There are two cases:

- (a) Every language $L \in \text{NP}$ either belongs to P or is NP-complete.
- (b) There exists an NP-intermediate language L .

Ladner's Theorem establishes that if $\text{P} \neq \text{NP}$, then there is an NP-intermediate language L [Lad75]. In light of Ladner's result, there has been a great deal of effort to find NP-intermediate problems. There have been numerous candidates since 1975. Some of these candidates have been shown to be in P. For others, their complexity status remains open. We provide some examples.

- (a) The Integer Factorization and Discrete Logarithm problems are conjectured to be NP-intermediate. To date, the best known algorithms for these problems rely on heavy-handed machinery from Algebraic Number Theory, such as Number Field Sieves.
- (b) Isomorphism problem such as Cayley Group Isomorphism and Graph Isomorphism are conjectured to be NP-intermediate. The trivial bound for Cayley Group Isomorphism is $n^{\log_p(n)+O(1)}$, where p is the smallest prime dividing n . This is obtained via a generator-enumeration technique, which Miller attributes to Tarjan [Mil78]. The best known algorithm for Cayley Group Isomorphism is $n^{(1/4)\log_p(n)+O(1)}$ due to Rosenbaum [Ros13] (see [LGR16, Sec. 2.2]). Even the impressive body of work on practical algorithms for this problem, led by Eick, Holt, Leedham-Green and O'Brien (e.g., [BEO02, ELGO02, BE99, CH03]) still results in an $n^{\Theta(\log n)}$ -time algorithm in the general case [Wil19].

For Graph Isomorphism, the best known algorithmic upper bound is $n^{\Theta(\log^2(n))}$ due to Babai [Bab15]. Cayley Group Isomorphism is many-to-one polynomial-time reducible to Graph Isomorphism, and so the Cayley Group Isomorphism problem remains a key bottleneck to determining whether Graph Isomorphism belongs to P.

- (c) Primality testing was a long-standing candidate to be NP-intermediate. In 2002, Agrawal, Kayal, and Saxena showed that Primality $\in \text{P}$ [AKS02]. It is noteworthy that Kayal and Saxena were both undergraduate students working with Agrawal. The techniques in their paper are largely accessible after a semester of Abstract Algebra at the undergraduate level. We also note that the implicit constant associated with the AKS procedure is quite large, which makes the procedure impractical. In practice, primality testing is still handled using probabilistic algorithms.
- (d) Linear Programming was long-conjectured to be NP-intermediate. The Simplex algorithm was the long-standing algorithmic tool in this area. Despite the fact that the Simplex algorithm performed well in practice, it still had a worst-case exponential runtime. In 1979, Leonid Khachiyan introduced the Ellipsoid algorithm, which was the first polynomial-time algorithm for Linear Programming [Kha80]. Thus, we have that Linear Programming $\in \text{P}$. It is worth noting that the implicit constant associated with the Ellipsoid algorithm is quite large. So in practice, the Simplex algorithm is still widely used today.

We now turn our attention to proving Ladner's Theorem. The weak version of Ladner's Theorem provides only one intermediate language, if $\text{P} \neq \text{NP}$.

Theorem 86 (Ladner (weak), 1975). Suppose that $\text{P} \neq \text{NP}$. Then there exists a language $L \in \text{NP}$ such that $L \notin \text{P}$ and L is not NP-complete.

Ladner proved something much stronger. Namely, if $P \neq NP$, then there is a strict infinite hierarchy of NP-intermediate languages. It is widely believed that $P \neq NP$. However, finding even one NP-intermediate language remains an open problem.

Theorem 87 (Ladner (strong), 1975). Suppose that $P \neq NP$. Then there exists a sequence of languages $(L_i)_{i \in \mathbb{N}}$ satisfying the following.

- (a) $L_{i+1} \subsetneq L_i$ for all $i \in \mathbb{N}$.
- (b) $L_i \notin P$ for all $i \in \mathbb{N}$.
- (c) L_i is not NP-complete for all $i \in \mathbb{N}$.
- (d) $L_i \not\leq_p L_{i+1}$ for all $i \in \mathbb{N}$.

In order to prove Theorems 86 and 87, we prove a more general theorem first.

Theorem 88. Suppose that $L \notin P$ is computable. Then there exists a language $K \notin P$ such that $K \leq_p L$ and $L \not\leq_p K$.

The key technique in proving Theorem 88 is to diagonalize against both polynomial-time Turing Machines to ensure that $K \notin P$ and polynomial-time reductions from L to K to ensure that $L \not\leq_p K$. We alternate between diagonalizing against polynomial-time Turing Machines at one stage, and then against polynomial-time reductions in the next stage.

Proof of Theorem 88. We construct the language K via a diagonalization argument. Namely, define:

$$K := \{x \in L : f(|x|) \text{ is even}\},$$

where $f(x)$ is a function we will construct later. Let $(M_i)_{i \in \mathbb{Z}^+}$ be an enumeration of polynomial-time deterministic Turing Machines, which enumerates P . Let $(F_i)_{i \in \mathbb{Z}^+}$ be an enumeration of polynomial-time Turing Machines without restriction to their output lengths. We note that $(F_i)_{i \in \mathbb{Z}^+}$ includes polynomial-time many-to-one reductions from L to K .

Let M_L be a decider for L . We now define f recursively as follows. First, define $f(0) = f(1) = 2$. We associate f with the Turing Machine M_f that computes f . On input 1^n (with $n > 1$), M_f proceeds in two stages, each lasting exactly n steps.⁵ At the first stage, M_f computes $f(0), f(1), \dots$, until it runs out of time. Suppose that the last value M_f computed at the first stage was $f(x) = k$. We will either have that $f(n) = k$ or $f(n) = k + 1$, depending on the second stage.

At the second stage, we have one of two cases.

- **Case 1:** Suppose that $k = 2i$. Here, we diagonalize against the i th language $L(M_i)$ in P as follows. The goal is to find a string z such that $z \in (L(M_i) \Delta K)$. We enumerate such strings z in lexicographic order, and then compute $M_i(z), M_L(z)$ ⁶, and $f(|z|)$ for all such strings. Note that by definition of K , we must compute $f(|z|)$ to ensure that $f(|z|)$ is even. If such a string is found in the allotted time (n steps), then M_f outputs $k + 1$ (so M_f can proceed to diagonalize against polynomial-time reductions on the next iteration). Otherwise, M_f outputs k (as we have not successfully diagonalized against M_i yet).
- **Case 2:** Suppose that $k = 2i - 1$. Here, we diagonalize against the i th polynomial-time computable function F_i . In this case, M_f searches for a string z such that F_i is not a valid Karp reduction on z . That is, either:
 - $z \in L$, but $F_i(z) \notin K$; or
 - $z \notin L$, but $F_i(z) \in K$.

⁵In order to prove Theorem 88, it is not necessary to limit the number of steps M_f can take to be polynomial in n . However, as $L \notin P$, M_L does not run in polynomial-time. So in order to apply Theorem 88 to prove Theorems 86 and 87, we need to build in the clocking condition into the proof here.

⁶By limiting M_f to n steps at the second stage, we ensure the (partial) computation of M_L does not run for so long that it prohibits K from belonging to NP.

We accomplish this by computing $F_i(z), M_L(z), M_L(F_i(z))$, and $f(|F_i(z)|)$. Here, we use clocking to ensure that M_L is not taking too long. If such a string z is found in the allotted time, then the output of M_f is $k + 1$. Otherwise, M_f outputs k .

We now show that K satisfies the following conditions.

- (a) $K \leq_p L$,
- (b) $K \notin \mathbf{P}$, and
- (c) $L \not\leq_p K$.

Proposition 89. $K \leq_p L$.

Proof. The inclusion map $\iota : K \rightarrow L$ sending $\iota(\omega) = \omega$ is a polynomial-time computable reduction. □

Proposition 90. $K \notin \mathbf{P}$.

Proof. Suppose to the contrary that $K \in \mathbf{P}$. Then there exists a polynomial time deterministic Turing Machine M_i such that $L(M_i) = K$. By Case 1 of the second stage in the construction on M_f , no string z was found satisfying $z \in K$ and $z \notin L(M_i)$ (or vice-versa). So $f(n)$ is even for all but finitely many n . So K and L coincide for all but finitely many strings. As $L \setminus K$ is finite, $L \setminus K$ can be decided in polynomial time. Together with the fact that $K \in \mathbf{P}$, we have that $L = K \cup (L \setminus K) \in \mathbf{P}$, contradicting the assumption that $L \notin \mathbf{P}$. □

Proposition 91. $L \not\leq_p K$.

Proof. Exercise. □

□

3.1.1 Exercises

(Recommended) Problem 31. Prove Proposition 91. [**Hint:** Argue by contradiction. Suppose that we have a polynomial-time computable reduction F_i that takes L to K . Argue that $f(n)$ will be even for only finitely many n .]

(Recommended) Problem 32. Using Theorem 88, prove Theorem 86.

(Recommended) Problem 33. Using Theorem 88, prove Theorem 87.

3.2 Introduction to Space Complexity

In this section, we introduce notions of space complexity. Our goal is to develop familiarity with the space complexity classes PSPACE, L, and NL, to the extent necessary to develop the theory of structural complexity. These complexity classes have a rich theory in and of themselves. We defer to [Sip96, AB09] for more comprehensive treatments.

In order to discuss space complexity, we consider a class of multitape Turing Machines, which we refer to as *space-bounded Turing Machines*. Precisely, we consider Turing Machines that use three tapes. These tapes are as follows:

- (a) The input tape, for which we may only read input.
- (b) The work tape, for which we have both read and write access.
- (c) The output tape, for which we have both read and write access.

Only non-blank cells on the work and output count towards the amount of space used. We now define the DSPACE and NSPACE complexity classes.

Definition 92. Let $f : \mathbb{N} \rightarrow \mathbb{N}$. We say that the language $L \in \text{DSPACE}(f(n))$ if there exists a deterministic space-bounded Turing Machine M such that M decides L and M uses at most $O(f(n))$ space. Similarly, we say that $L \in \text{NSPACE}(f(n))$ if there exists a non-deterministic space-bounded Turing Machine M such that M decides L and M uses at most $O(f(n))$ space.

Clearly, $\text{DSPACE}(f(n)) \subseteq \text{NSPACE}(f(n))$. We next relate our space and time complexity classes.

Lemma 93. We have that $\text{DTIME}(f(n)) \subseteq \text{DSPACE}(f(n))$.

Proof. We note that a deterministic Turing Machine that takes k steps writes to at most k cells. The result follows. \square

We now show that $\text{NSPACE}(f(n)) \subseteq \text{DTIME}(2^{O(f(n))})$.

Theorem 94. We have that $\text{NSPACE}(f(n)) \subseteq \text{DTIME}(2^{O(f(n))})$.

Proof. The idea is to count Turing Machine configurations. Each Turing Machine state transition costs a single computation step and takes us from one configuration to another. Furthermore, each configuration is visited at most once; otherwise the Turing Machine enters an infinite loop⁷. So counting the number of Turing Machine configurations provides an upper bound on the runtime.

Let $L \in \text{NSPACE}(f(n))$, and suppose that we have a non-deterministic space-bounded Turing Machine M that accepts L . Without loss of generality, suppose that M uses exactly $f(n)$ tape cells. Each configuration consists of the following:

- (a) A state from $Q(M)$.
- (b) The positions of the input, work, and output tape heads. There are n possible positions for the input tape head, at most $f(n)$ possible positions for the work tape head, and at most $f(n)$ possible positions for the output tape head. The selections of the positions for each tape head are independent. So by the rule of product, there are at most $n \cdot (f(n))^2$ possible selections for the tape head positions.
- (c) There are at most $f(n)$ non-blank cells between the work and output tapes. So there are at most $|\Gamma|^{f(n)}$ possible ways to fill the non-blank cells.

⁷Implicitly, we are describing a directed acyclic graph structure, which is commonly referred to as a *configuration graph*. We are not leveraging the graph structure in this proof, so we omit those details. However, the configuration graph is a useful tool in space complexity, such as in establishing a natural NL-complete problem.

Note that selecting the state, positions for the tape heads, and ways to fill the non-blank cells are independent selections. So by the rule of product, we have at most:

$$|Q(M)| \cdot n \cdot (f(n))^2 \cdot |\Gamma|^{f(n)} \leq |\Gamma|^{O(f(n))}$$

possible configurations. Now note that $k = 2^{\log_2(k)}$. So if $k := |\Gamma| > 2$, we have that:

$$\begin{aligned} |\Gamma|^{O(f(n))} &= (2^{\log_2(k)})^{O(f(n))} \\ &= 2^{\log_2(k) \cdot O(f(n))} \\ &= 2^{O(f(n))}. \end{aligned}$$

The result follows. □

We conclude by introducing some standard complexity classes and complete problems.

3.2.1 PSPACE

Definition 95. The complexity class PSPACE is defined to be:

$$\text{PSPACE} := \bigcup_{k \in \mathbb{N}} \text{DSPACE}(n^k).$$

PSPACE also has complete problems under polynomial-time reductions. We introduce one such canonical problem here- TQBF, though we won't prove that $\text{TQBF} \in \text{PSPACE}$.

Definition 96. Let $\varphi(x_1, \dots, x_n)$ be a Boolean formula. We consider *fully quantified Boolean formulas*, which are of the form:

$$\Phi := Q_1 x_1, Q_2 x_2, \dots, Q_n x_n \varphi(x_1, \dots, x_n).$$

Here, each quantifier Q_i is either an existential quantifier (\exists) or a universal quantifier (\forall). We note that a fully quantified Boolean formula is either always true or always false.

Definition 97 (TQBF). The True Quantified Boolean Formulas (TQBF) problem is defined as follows:

- Instance: A fully quantified Boolean formula Φ .
- Decision: Is Φ true?

Example 98. Let:

$$\Phi := \forall x \exists y [(x \vee y) \wedge (\bar{x} \vee \bar{y})].$$

Observe that Φ is true: take $y = \bar{y}$. On the other hand,

$$\Psi := \exists x \forall y [(x \vee y) \wedge (\bar{x} \vee \bar{y})]$$

is false. If $x = y$, then either $(x \vee y)$ is true or $(\bar{x} \vee \bar{y})$ is true (but not both).

Theorem 99. We have that TQBF is PSPACE-complete.

Remark 100. We note that TQBF remains PSPACE-complete if we insist that the Boolean formula is in conjunctive normal form, where each clause has exactly 3 literals (that is, the Boolean formula is in 3 – CNF). We will use this fact later when discussing Interactive Proofs.

3.2.2 L and NL

Definition 101. The complexity class $L := \text{DSPACE}(\log(n))$. Similarly, $NL := \text{NSPACE}(\log(n))$.

We note that NL has complete problems, under logspace reductions. We denote the relation that L_1 is logspace many-to-one reducible to L_2 as $L_1 \leq_\ell L_2$. Informally, complete problems capture the essence of a given complexity class. For this reason, it is natural to insist that reductions are computable within said complexity class. Hence, we consider complete problems for NL under logspace reductions.

Remark 102. We note that every non-trivial problem in L is L-complete under logspace reductions. For this reason, complete problems are often considered under more restrictive notions of reducibility, such as AC^0 -reducibility. We won't pursue this direction any further.

The main result we consider in this section is in establishing an NL-complete problem. Namely, the Connectivity problem.

Definition 103. The Connectivity problem is defined as follows.

- Instance: A graph $G(V, E)$, and vertices u and v .
- Decision: Does there exist a $u \rightarrow v$ path in G ?

Theorem 104. The Connectivity problem is NL-complete.

Our proof is adopted from [Kat11].

Proposition 105. We have that $\text{Connectivity} \in \text{NL}$.

Proof. Let $G(V, E)$ be our graph on n vertices, and let $u, v \in V(G)$ be the specified start and end vertices. We exhibit a non-deterministic algorithm that always rejects if no such $u \rightarrow v$ path exists and sometimes accepts if there exists a $u \rightarrow v$ path. The idea is to non-deterministically guess the vertices in the path one at a time. The algorithm proceeds as follows.

1. If $u = v$, we terminate.
2. Set $v_{\text{current}} := u$.
3. For $i = 0$ to n :
 - (a) Guess a vertex w .
 - (b) If w is not neighbor of v_{current} , then reject.
 - (c) If $w = v$, accept.
 - (d) Otherwise, set $v_{\text{current}} = w$.
4. If we have not decided after the loop terminates, then we reject, as any path has length at most n .

We note that the vertices are specified by binary strings of length $\lceil \log_2(n) \rceil$. At any given iteration, we store v_{current} and our guess w . So only $2\lceil \log_2(n) \rceil \in O(\log(n))$ bits of memory are being stored. Thus, $\text{Connectivity} \in \text{NL}$. \square

We now show that Connectivity is NL-hard. The idea is to take an NL-TM and map it to its corresponding configuration graph. We then ask whether there is a path from the initial configuration to a specified accepting configuration.

Proposition 106. We have that Connectivity is NL-hard.

Proof. Let $L \in \text{NL}$, and let M be an NL-TM that accepts L . Fix a string $\omega \in \Sigma^*$ of length n . We note that as M , when run on ω , uses space $c \cdot \log(n)$, then M has $2^{c \cdot \log(n)} = n^c$ possible configurations. We construct a configuration graph $G(V, E)$, where the configurations are the vertices. Two configurations C_1 and C_2 are adjacent if there is a state transition of M that takes us from C_1 to C_2 . Let C_0 be the initial configuration. By construction, we have that $\omega \in L$ if and only if there exists an accepting configuration C^* that is reachable from C_0 .

We note that we don't construct the configuration graph outright, as this would require at most $O(n^{2c})$ cells in memory to store the graph. Instead, we note that at any step, we only compute the configurations corresponding to the initial and final states, as well as the given configuration we are considering at any iteration when we are guessing the path. As there are n^c configurations, only $c \cdot \lceil \log_2(n) \rceil$ bits are required to store a given configuration. We are storing $3c \cdot \lceil \log_2(n) \rceil \in O(\log(n))$ bits in memory at a given stage. So our reduction is logspace computable, as desired. \square

3.2.3 Exercises

(Recommended) Problem 34. Theorem 94 yields some immediate corollaries.

- (a) Show that a logspace computation can be simulated in polynomial time. That is, using Theorem 94 show that: $\text{NL} \subseteq \text{P}$. [**Note:** As a result, we obtain that logspace-computable many-to-one reductions are stronger than polynomial-time computable many-to-one reductions. That is, $L_1 \leq_\ell L_2 \implies L_1 \leq_p L_2$.]
- (b) Define:

$$\text{EXPTIME} = \bigcup_{k \in \mathbb{N}} \text{DTIME}(2^{O(n^k)}).$$

Show that $\text{PSPACE} \subseteq \text{EXPTIME}$.

(Recommended) Problem 35. A *finite state automaton* is a five-tuple $(Q, \Sigma, \delta, q_0, F)$, where:

- Q is our finite set of states,
- Σ is our alphabet,
- $\delta : Q \times \Sigma \rightarrow Q$ is our transition function,
- q_0 is the initial state, and
- $F \subseteq Q$ is the set of accepting states.

A finite state automaton parses the input string $\omega = (\omega_1, \dots, \omega_n)$ one character at a time in order, starting from ω_1 . No characters are revisited, and the finite state automaton does not have a work tape. The finite state automaton writes either 0 or 1 to the output tape, depending on whether it accepts the given string. Do the following.

- (a) Design a finite state automaton to accept the language $L \subseteq \{0, 1\}^*$, where each string in L has an even number of 1's.
- (b) The set of languages accepted by finite state automata are precisely the regular languages, denoted REG . Show that $\text{REG} = \text{DSPACE}(O(1))$. [**Hint:** To show that $\text{DSPACE}(O(1)) \subseteq \text{REG}$, show that the configuration graph for a $\text{DSPACE}(O(1))$ -TM can be viewed as a deterministic finite state automaton. You may assume, without loss of generality, the TMs may have multiple accept states.]
- (c) A *non-deterministic finite state automaton* is defined similarly as a deterministic finite state automaton $(Q, \Sigma, \delta, q_0, F)$, except where the transition function is of the form:

$$\delta : Q \times \Sigma \rightarrow 2^Q.$$

Kleene's Theorem provides that L is accepted by some non-deterministic finite state automaton if and only if $L \in \text{REG}$.⁸ Show that $\text{REG} = \text{NSPACE}(O(1))$.

- (d) Conclude that $\text{DSPACE}(O(1)) = \text{NSPACE}(O(1))$.

⁸More precisely, Kleene's Theorem states that $L \in \text{REG}$ if and only if L is accepted by some deterministic finite state automaton. The need to consider non-deterministic variants comes out in the proof, where it is subsequently shown that non-deterministic finite state automata can be converted to equivalent deterministic finite state automata. We defer to [Lev20] for a more thorough treatment of regular languages.

(Recommended) Problem 36. Define:

$$\text{NPSpace} := \bigcup_{k \in \mathbb{N}} \text{NSpace}(n^k).$$

Show that $\text{PSPACE} = \text{NPSpace}$. [**Hint:** Suppose that an NPSpace machine uses space $O(n^k)$. Then there are $2^{O(n^k)}$ possible configurations. How much space does it cost to store a single configuration? What about to traverse the configuration graph?]

(Recommended) Problem 37. While we have shown that $\text{P} \subseteq \text{PSPACE}$ (see Lemma 93), we have not shown that $\text{NP} \subseteq \text{PSPACE}$. To do so, show that $\text{SAT} \in \text{PSPACE}$. [**Hint:** Try a brute force approach to solve SAT.]

(Recommended) Problem 38. Show that $\text{TQBF} \in \text{PSPACE}$. [**Hint:** Apply the same technique as in Problem 37.]

3.3 Baker-Gill-Solovay Theorem

Computational Complexity began as a generalization of Computability Theory. The key idea is to ask what problems can and cannot be solved under given resource constraints. Many of the core techniques in Structural Complexity Theory, such as diagonalization and simulation, arise from Computability Theory. As a result, it is very natural to ask whether Computability techniques are sufficient to resolve the P vs. NP problem. In 1975, Baker, Gill, and Solovay showed that Computability techniques are insufficient to resolve the P vs. NP problem. The key observation is that Computability techniques are not sensitive to the presence of oracles. However, there are oracles A and B such that $P^A = NP^A$, but $P^B \neq NP^B$.

Theorem 107 (Baker-Gill-Solovay, 1975.). There exist oracles A and B such that $P^A = NP^A$, but $P^B \neq NP^B$.

We begin with the oracle A . It suffices to take A to be a PSPACE-complete problem.

Proposition 108. Let A be PSPACE-complete. Then $P^A = NP^A$.

We note that $P^A \subseteq NP^A$ for the same reasons that $P \subseteq NP$. The proof that $P \subseteq NP$ is not sensitive to the presence of oracles. To show that $NP^A \subseteq P^A$, it suffices to show that the following chain of inclusions holds:

$$NP^A \subseteq PSPACE \subseteq P^A.$$

We leave it as an exercise to establish this chain of inclusions. See Problem 39.

We now turn towards constructing an oracle B , for which $P^B \neq NP^B$.

Proposition 109. There exists an oracle B such that $P^B \neq NP^B$.

Proof. We construct an oracle B inductively; such that for each $i \in \mathbb{N}$, B contains at most one word of length i . Precisely, we construct a sequence of oracles $(B_i)_{i \in \mathbb{N}}$ such that:

$$B := \bigcup_{i \in \mathbb{N}} B_i.$$

Our strategy is then to show that the following language:

$$L := \{0^i : B \text{ has a string of length } i\}$$

is in $NP^B \setminus P^B$. In the process of constructing B , we also maintain a list of forbidden words. Denote \mathcal{F}_i to be the set of forbidden words encountered after constructing B_i .

We proceed via diagonalization, constructing oracles $(B_i)_{i \in \mathbb{N}}$. Let $(M_\ell)_{\ell \in \mathbb{N}}$ be an enumeration of Oracle Turing Machines. When we consider M_ℓ , we will equip M_ℓ with the oracle B_ℓ (that is, we consider $M_\ell^{B_\ell}$).⁹ Define $B_0 := \emptyset$ and $\mathcal{F}_0 := \emptyset$. Fix $i \geq 0$, and suppose that we have constructed B_i , where B_i has at most one word of length j for each $0 \leq j < i$, and no words of length at least i . Similarly, suppose that we have constructed the set \mathcal{F}_i of forbidden words. We simulate $M_i^{B_i}$ on 0^i for $i^{\log i}$ steps. If $M_i^{B_i}$ queries a word y of length at least i , we assume that $y \notin B$ and add y to the set of forbidden words. Let:

$$\mathcal{F}_{i+1} := \mathcal{F}_i \cup \{y \in \Sigma^* : |y| \geq i \text{ and } M_i^{B_i} \text{ queried } y\}.$$

We now turn our attention to constructing B_{i+1} . We have the following cases:

- **Case 1:** Suppose that in the first $i^{\log i}$ steps that $M_i^{B_i}$ rejects 0^i . If there is a word ω of length i that does not belong to \mathcal{F}_i , then set $B_{i+1} := B_i \cup \{\omega\}$. Such a word ω may be selected deterministically, such as by selecting the first non-forbidden word in lexicographic order of Σ^i . If no such word ω exists, then $B_{i+1} := B_i$.
- **Case 2:** Suppose that $M_i^{B_i}$ does not reject 0^i in the first $i^{\log i}$ steps. Then no word of length i is placed in B .

⁹Recall that an M_ℓ is defined by its transition function. The oracle is a parameter, and not part of the definition of M_ℓ .

We also do not include a word of length i in B_{i+1} if all words of length i are forbidden. Note that for sufficiently large i , not all words of length i will be forbidden. To see this, observe that the maximum number of forbidden words of length at most i is:

$$\sum_{j=1}^i j^{\log j} \leq i(i^{\log i}) = i^{1+\log i}.$$

Now not all words of length i are forbidden if $|\Sigma|^i \geq 2^i > i^{1+\log i}$. Note that $2^i > i^{1+\log i}$ holds precisely if $i > (\log i)(1 + \log i)$. This last inequality holds for $i \geq 32$. In particular, it follows that B is an infinite set. We again note that:

$$B := \bigcup_{i \in \mathbb{N}} B_i.$$

and

$$L := \{0^i : B \text{ has a word of length } i\}.$$

We claim that $L \in \text{NP}^B \setminus \text{P}^B$.

Lemma 110. We have that $L \in \text{NP}^B$.

Proof. Exercise. □

Lemma 111. We have that $L \notin \text{P}^B$.

Proof. Suppose to the contrary that $L \in \text{P}^B$. So there exists an Oracle Turing Machine M_k such that $L = L(M_k^B)$. Let $p(n)$ be the runtime complexity function for M_k^B . As B is infinite, so is L . Thus, M_k accepts arbitrarily long strings. So we may- without loss of generality- assume that $k \geq 32$ and $2^k \geq p(k)$. We now analyze whether $0^k \in L$. We have the following cases.

- **Case 1:** Suppose that M_k^B accepts 0^k . So $0^k \in L$, which implies that B has a string ω of length k . By construction of B , it follows that $M_k^{B_k}$ rejects 0^k . However, as $B_k \subseteq B$ and B has no words of length at least k that are queried by $M_k^{B_k}$ on 0^k , it follows that $M_k^{B_k}$ and M_k^B behave identically on 0^k . So M_k^B rejects 0^k , a contradiction.
- **Case 2:** Suppose that M_k^B rejects 0^k . We leave it as an exercise to derive a contradiction by showing that this implies that $0^k \in L$. See Problem 41. □

3.3.1 Exercises

(Recommended) Problem 39. Let A be PSPACE-complete. Do the following.

- (a) Show that $\text{NP}^A \subseteq \text{PSPACE}$.
- (b) Show that $\text{PSPACE} \subseteq \text{P}^A$.
- (c) Conclude that $\text{P}^A = \text{NP}^A$.

(Recommended) Problem 40. Prove Lemma 110.

(Recommended) Problem 41. Complete the proof of Lemma 111. Suppose that M_k^B rejects 0^k . Derive a contradiction by showing that this implies that $0^k \in L$. □

3.4 Polynomial-Time Hierarchy: Introduction

We recall the definitions of NP and coNP.

Definition 112 (NP). We say that a language $L \in \text{NP}$ if there exists a polynomial $p(\cdot)$ depending only on L and a verifier V such that if $x \in L$, then there exists a string y of length at most $p(|x|)$, such that $V(x, y) = 1$ and V runs in time $O(p(|x|))$. Here, y is the *certificate*.

Remark 113. We may write the definition of NP in quantifier notation:

$$x \in L \iff \exists y \text{ s.t. } |y| \leq p(|x|), V(x, y) = 1. \quad (11)$$

This expression is often abbreviated as follows, though (11) is the formalism and intended meaning.

$$x \in L \iff \exists y, V(x, y) = 1.$$

Similarly, coNP is defined as follows.

Definition 114 (coNP). We say that a language $L \in \text{coNP}$ if there exists a polynomial $p(\cdot)$ depending only on L and a verifier V such that if $x \in L$, then for every y of length at most $p(|x|)$, that $V(x, y) = 0$ and V runs in time $O(p(|x|))$.

Remark 115. We may write the definition of coNP in quantifier notation:

$$x \in L \iff \forall y \text{ s.t. } |y| \leq p(|x|), V(x, y) = 0. \quad (12)$$

This expression is often abbreviated as follows, though (12) is the formalism and intended meaning.

$$x \in L \iff \forall y, V(x, y) = 0.$$

Our goal now is to generalize NP and coNP. We begin with some motivation. Recall the **Independent Set** decision problem, which takes as input a graph $G(V, E)$ and integer k , and asks if G has an independent set of size k . Recall that **Independent Set** is NP-Complete. In particular, **Independent Set** \in NP.

Now consider the **Maximum Independent Set** problem, which again takes as input a graph $G(V, E)$ and integer k . Here, we ask whether the largest size independent set in G has k vertices. Here, we need to verify a couple conditions:

- G has an independent set of k vertices. This is precisely the condition that **Independent Set** \in NP.
- G does not have an independent set of $k + 1$ vertices. We note that verifying this second condition is a coNP problem.

So effectively, $(G, k) \in \text{Maximum Independent Set} \iff$ there exists a small certificate of one type and no small certificate of another type. This motivates the definition of a new complexity class, which we will call Σ_2^P .

Definition 116. We say that a language $L \in \Sigma_2^P$ (pronounced Sigma-2) if there exists a polynomial $p(\cdot)$ depending only on L and a verifier V such that if $\omega \in L$, then there exists a string x of length at most $p(|\omega|)$, such that for all strings y of length at most $p(|\omega|)$, $V(\omega, x, y) = 1$ and V runs in time $O(p(|\omega|))$.

We may again express the definition of Σ_2^P in quantifier notation.

$$\omega \in L \iff \exists x \text{ s.t. } |x| \leq p(|\omega|), \forall y \text{ s.t. } |y| \leq p(|\omega|), M(\omega, x, y) = 1.$$

As we saw with NP and coNP, the quantified expression for Σ_2^P is commonly abbreviated as follows.

$$\omega \in L \iff \exists x, \forall y, M(\omega, x, y) = 1.$$

Example 117. We note that **Maximum Independent Set** $\in \Sigma_2^P$. Here, x is the certificate for the independent set of size k , and y is a vertex set of size $k + 1$. In other words, M checks that x is an independent set of size k and that y is a $(k + 1)$ -size vertex set is not an independent set. Note that M itself does not consider all such $(k + 1)$ -size vertex sets at once. Rather, the quantifier states that for any given $(k + 1)$ -size vertex set y , that M will check that y is not an independent set.

3.4.1 Σ_i^p

We now turn to generalizing Σ_2^p . Here, the subscript 2 indicates that we use two quantifiers. We define Σ_i^p to use i quantifiers, starting with \exists , and then alternating between \exists and \forall . This is formalized as follows.

Definition 118. We say that the language $L \in \Sigma_i^p$ if there exists a polynomial $p(\cdot)$ depending only on L and a verifier V such that:

$$\omega \in L \iff \exists x_1, \forall x_2, \exists x_3, \forall x_4, \dots, Q_i x_i, V(\omega, x_1, \dots, x_i) = 1,$$

$|x_j| \leq p(|\omega|)$ for all $1 \leq j \leq i$, and V runs in time $O(p(|\omega|))$. Note that Q_i indicates a quantifier. In particular, Q_i is an existential quantifier if i is odd and a universal quantifier if i is even.

So for Σ_3^p , the abbreviated quantified expression is:

$$\omega \in L \iff \exists x_1, \forall x_2, \exists x_3, V(\omega, x_1, x_2, x_3) = 1.$$

Similarly, for Σ_4^p , the abbreviated quantified expression is:

$$\omega \in L \iff \exists x_1, \forall x_2, \exists x_3, \forall x_4, V(\omega, x_1, x_2, x_3, x_4) = 1.$$

Remark 119. It is also worth noting that $\text{NP} = \Sigma_1^p$.

Theorem 120. Fix $i \geq 1$. We have that $\Sigma_i^p \subseteq \Sigma_{i+1}^p$.

Proof. Let $L \in \Sigma_i^p$. Let $M(\omega; x_1, \dots, x_i)$ be the Σ_i^p Turing Machine that accepts L . We construct a Σ_{i+1}^p Turing Machine $N(\omega; x_1, \dots, x_i, x_{i+1})$ that accepts L as follows. On input $(\omega; x_1, \dots, x_i, x_{i+1})$, N runs M on $(\omega; x_1, \dots, x_i)$. N accepts ω if and only if $M(\omega; x_1, \dots, x_i) = 1$. As M runs in polynomial time, so does N . The result follows. \square

3.4.2 Π_i^p

We now turn our attention to generalizing coNP . Note that the quantified expression for coNP is obtained by negating the quantifiers for NP . We define the complexity class $\Pi_i^p := \text{co}\Sigma_i^p$. That is, we negate the quantified expression for Σ_i^p to obtain a similar definition regarding alternating quantifiers. However, we begin with a universal quantifier rather than an existential quantifier. This definition is formalized as follows.

Definition 121. We say that the language $L \in \Pi_i^p$ if there exists a polynomial $p(\cdot)$ depending only on L and a verifier V such that if $\omega \in L$

$$\omega \in L \iff \forall x_1, \exists x_2, \forall x_3, \exists x_4, \dots, Q_i x_i, V(\omega, x_1, \dots, x_i) = 0,$$

$|x_j| \leq p(|\omega|)$ for all $1 \leq j \leq i$, and V runs in time $O(p(|\omega|))$. Note that Q_i indicates a quantifier. In particular, Q_i is an existential quantifier if i is even and a universal quantifier if i is odd.

So for Π_3^p , the abbreviated quantified expression is:

$$\omega \in L \iff \forall x_1, \exists x_2, \forall x_3, V(\omega, x_1, x_2, x_3) = 0.$$

Similarly, for Π_4^p , the abbreviated quantified expression is:

$$\omega \in L \iff \forall x_1, \exists x_2, \forall x_3, \exists x_4, V(\omega, x_1, x_2, x_3) = 0.$$

We now establish the following relations amongst the classes of the polynomial time hierarchy.

Theorem 122. Show that $\Pi_i^p \subseteq \Pi_{i+1}^p$.

Proof. Exercise \square

Theorem 123. Show that $\Sigma_i^p \subseteq \Pi_{i+1}^p$.

Proof. Exercise. \square

Theorem 124. Show that $\Pi_i^p \subseteq \Sigma_{i+1}^p$.

Proof. Exercise. □

As a final note, we define the Polynomial-Time Hierarchy formally:

Definition 125. The *Polynomial-Time Hierarchy*, denoted PH, is:

$$\text{PH} = \bigcup_{i \in \mathbb{N}} \Sigma_i^p.$$

Remark 126. It is conjectured that $\Sigma_i^p \neq \Pi_i^p$ for each $i \geq 1$. In particular, this conjecture implies that Σ_i^p and Π_i^p are both strictly contained in Σ_{i+1}^p and Π_{i+1}^p .

3.4.3 Exercises

(Recommended) Problem 42. Prove Theorem 122. Show that $\Pi_i^p \subseteq \Pi_{i+1}^p$.

(Recommended) Problem 43. Prove Theorem 123. Show that $\Sigma_i^p \subseteq \Pi_{i+1}^p$.

(Recommended) Problem 44. Prove Theorem 124. Show that $\Pi_i^p \subseteq \Sigma_{i+1}^p$.

3.5 Structure of Polynomial-Time Hierarchy

It is conjectured that $\Sigma_i^p \neq \Pi_i^p$ for each $i \geq 1$. In particular, this conjecture implies that Σ_i^p and Π_i^p are both strictly contained in Σ_{i+1}^p and Π_{i+1}^p . In this section, we examine conditions under which PH collapses, as well as consequences thereof.

Theorem 127. Suppose that for some $i \geq 1$, we have that $\Sigma_i^p = \Pi_i^p$. Then $\text{PH} = \Sigma_i^p$.

In order to prove Theorem 127, we first show that if $\Sigma_i^p = \Pi_i^p$, then $\Sigma_i^p = \Sigma_{i+1}^p = \Pi_{i+1}^p$. Theorem 127 then follows by induction.

Proposition 128. Suppose that for some $i \geq 1$, we have that $\Sigma_i^p = \Pi_i^p$. Then $\Sigma_i^p = \Sigma_{i+1}^p$.

Proof. By Theorem 120, we have that $\Sigma_i^p \subseteq \Sigma_{i+1}^p$. We show that $\Sigma_{i+1}^p \subseteq \Sigma_i^p$. Let $L \in \Sigma_{i+1}^p$. So there exists a Σ_{i+1}^p Turing Machine $M(\omega; x_1, \dots, x_{i+1})$ that accepts L . Suppose that $\omega \in L$. Then the following relation holds:

$$\exists x_1, \forall x_2, \dots, Qx_{i+1} M(\omega; x_1, \dots, x_{i+1}) = 1.$$

Define:

$$L' = \{\langle \omega, x_1 \rangle : \forall x_2, \dots, Qx_{i+1} M(\omega; x_1, \dots, x_{i+1}) = 1\}.$$

So M is a Π_i^p Turing Machine that accepts L' . Note that $L \in \Sigma_{i+1}^p$ if and only if $L' \in \Pi_i^p$. As $\Pi_i^p = \Sigma_i^p$, there exists a Σ_i^p Turing Machine $N(\langle \omega, x_1; x_2, \dots, x_{i+1} \rangle)$ that accepts L' . We define a Σ_i^p Turing Machine $N(\omega; \langle x_1, x_2 \rangle, \dots, x_{i+1})$ that accepts L as follows. On input $(\omega; \langle x_1, x_2 \rangle, x_3, \dots, x_{i+1})$, N' simulates $N(\langle \omega, x_1 \rangle; x_2, \dots, x_{i+1})$. Now N' accepts ω if and only if N accepts ω . So $L \in \Sigma_i^p$, as desired. \square

It remains to show that if $\Sigma_i^p = \Pi_i^p$, then $\Sigma_i^p = \Pi_{i+1}^p$. The key idea here is to use the fact that $\Sigma_i^p = \Pi_i^p = \Sigma_{i+1}^p$, by the previous proposition, as well as the fact that $\Pi_{i+1}^p = \text{co}\Sigma_{i+1}^p$. We leave this as an exercise.

Proposition 129. Suppose that for some $i \geq 1$, we have that $\Sigma_i^p = \Pi_i^p$. Then $\Sigma_i^p = \Pi_{i+1}^p$.

We similarly obtain that if $\text{P} = \text{NP}$, then $\text{PH} = \text{P}$. The key observation is that if $\text{P} = \text{NP}$, then $\text{NP} = \text{coNP}$. Recall that $\Sigma_1^p = \text{NP}$ and $\Pi_1^p = \text{coNP}$. Theorem 127 now applies immediately. We leave it as an exercise to fill in the details.

Theorem 130. Suppose that $\text{P} = \text{NP}$. Then $\text{PH} = \text{P}$.

We next discuss the relationship between PH and PSPACE. We first observe that if PH has a complete problem, then PH collapses to some level i .

Theorem 131. Suppose that PH has a complete problem under polynomial-time reductions. Then $\text{PH} = \Sigma_i^p$ for some $i \geq 1$.

As PSPACE has complete problems under polynomial-time reductions, we note that if $\text{PH} = \text{PSPACE}$, then PH collapses.

Remark 132. While PH is unlikely to have complete problems, both Σ_i^p and Π_i^p have complete problems. Establishing these complete problems is quite involved, similar to proving the Cook-Levin Theorem. We omit the proofs for this reason.

Definition 133 ($\Sigma_i^p\text{SAT}$).

- Instance: Let φ be a Boolean formula.
- Decision: Does the following relation hold?

$$\exists x_1 \forall x_2, \dots, Qx_i \varphi(x_1, \dots, x_i) = 1,$$

where $x_1, \dots, x_k \in \{0, 1\}$.

Definition 134 ($\Pi_i^p\text{SAT}$).

- Instance: Let φ be a Boolean formula.
- Decision: Does the following relation hold?

$$\forall x_1 \exists x_2, \dots, Qx_i \varphi(x_1, \dots, x_i) = 1,$$

where $x_1, \dots, x_k \in \{0, 1\}$.

Theorem 135. For each $i \geq 1$, $\Sigma_i^p\text{SAT}$ is Σ_i^p -complete, and $\Pi_i^p\text{SAT}$ is Π_i^p -complete.

3.5.1 Exercises

(Recommended) Problem 45. Prove Proposition 129. Suppose that for some $i \geq 1$, we have that $\Sigma_i^p = \Pi_i^p$. Then $\Sigma_i^p = \Pi_{i+1}^p$.

(Recommended) Problem 46. The goal of this problem is to prove Theorem 130. That is, if $P = NP$, then $PH = P$.

- (a) Show that if $P = NP$, then $NP = \text{coNP}$.
- (b) Apply Theorem 127 to show that if $P = NP$, then $P = PH$.

(Recommended) Problem 47. The goal of this problem is to show that it is unlikely that $PH = \text{PSPACE}$.

- (a) Prove Theorem 131. Suppose that PH has a complete problem under polynomial-time reductions. Then $PH = \Sigma_i^p$ for some $i \geq 1$.
- (b) Show that if PH does not collapse, then $PH \neq \text{PSPACE}$.

(Recommended) Problem 48. Show that $PH \subseteq \text{PSPACE}$ in two ways.

- (a) First, give a simulation argument. Show how to design a PSPACE algorithm for $\Sigma_i^p\text{SAT}$ (or if you prefer, for $\Pi_i^p\text{SAT}$).
- (b) Second, reduce to TQBF. [**Hint:** Observe that $\Sigma_i^p\text{SAT}$ and $\Pi_i^p\text{SAT}$ instances are also instances of TQBF.]

3.6 Polynomial-Time Hierarchy and Oracles

We may alternatively define the complexity classes in PH in terms of oracles. This is more in the spirit of the Arithmetic Hierarchy. Namely, denote $\Sigma_0^P = \Pi_0^P = P$. For $i \geq 1$, we have that $\Sigma_i^P = \text{NP}^{\Sigma_{i-1}^P}$ and $\Pi_i^P = \text{coNP}^{\Pi_{i-1}^P}$. We prove the former relation here. The latter falls out as a corollary, which we leave as an exercise.

Theorem 136. For each $i \geq 1$, we have that $\Sigma_i^P = \text{NP}^{\Sigma_{i-1}^P}$.

Proof. The proof is by induction on i . When $i = 1$, we have that:

$$\begin{aligned}\Sigma_1^P &= \text{NP}^{\Sigma_0^P} \\ &= \text{NP}^P \\ &= \text{NP}.\end{aligned}$$

Fix $i \geq 1$, and suppose that $\Sigma_i^P = \text{NP}^{\Sigma_{i-1}^P}$. We now show that $\Sigma_{i+1}^P = \text{NP}^{\Sigma_i^P}$.

Lemma 137. We have that $\Sigma_{i+1}^P \subseteq \text{NP}^{\Sigma_i^P}$.

Proof. Exercise. □

Lemma 138. We have that $\text{NP}^{\Sigma_i^P} \subseteq \Sigma_{i+1}^P$.

Proof. Let $L \in \text{NP}^{\Sigma_i^P}$, and let $M^{\Sigma_i^P}$ be the corresponding Oracle TM. Let $\omega \in L$. Suppose that y is a corresponding certificate, and that $M^{\Sigma_i^P}$ makes m queries: $(q_1, a_1), \dots, (q_m, a_m)$, where the q_i are the queries and the a_i are the answers. In particular, the following are equivalent:

- (a) $\omega \in L$.
- (b) There exists a certificate y such that $M^{\Sigma_i^P}(\omega, y) = 1$.
- (c) There exist queries $(q_1, a_1), \dots, (q_m, a_m)$, where both m and the (q_j, a_j) pairs depend on y , such that:

$$\bigwedge_{j=1}^m M_j(q_j) = a_j,$$

for Σ_i^P Turing Machines M_1, \dots, M_m .

We note that if $M_j(q_j) = 1$, then $q_j \in L(M_j)$. In particular, as M_j is a Σ_i^P Turing Machine, we have by the inductive hypothesis that:

$$\exists y_{j,1}, \forall y_{j,2}, \dots, \exists y_{j,i} R_j(q_j, y_{j,1}, \dots, y_{j,i}) = 1, \quad (13)$$

where R_j is the Σ_i^P relation corresponding to $L(M_j)$. Now if $M_j(q_j) = 0$, then $q_j \notin L(M_j)$. Now $\overline{L(M_j)} \in \Pi_i^P$. It follows that:

$$\forall y'_{j,2}, \exists y'_{j,3}, \dots, \exists y'_{j,i+1} R_j(q_j, y'_{j,2}, \dots, y'_{j,i+1}) = 0. \quad (14)$$

Combining (13) and (14), we have that $M_j(q_j) = a_j$ if and only if:

$$\begin{aligned}&\exists y_{j,1}, \forall (y_{j,2}, y'_{j,2}), \dots, \exists y_{j,i}, \forall (y'_{j,i}, y_{j,i+1}), \\ &[a_j = 1 \wedge R_j(q_j, y_{j,1}, \dots, y_{j,i}) = 1] \vee [a_j = 0 \wedge R_j(q_j, y'_{j,2}, \dots, y'_{j,i+1}) = 0].\end{aligned}$$

Our strategy now is to collect the $y_{j,k}$ and $y'_{j,k}$ strings into a Σ_{i+1}^P formula. We have that $M(\omega) = 1$ if and only if:

$$\exists q_1, \dots, q_m, y, y_{1,1}, \dots, y_{m,1}, \quad (15)$$

$$\forall y_{1,2}, \dots, y_{m,2}, y'_{1,2}, \dots, y'_{m,2}, \quad (16)$$

$$\exists y_{1,3}, \dots, y_{m,3}, y'_{1,3}, \dots, y'_{m,3}, \quad (17)$$

$$\vdots \quad (18)$$

$$Q_i(y_{1,i}, \dots, y_{m,i}, y'_{1,i}, \dots, y'_{m,i}), \quad (19)$$

$$Q_{i+1}(y'_{1,i+1}, \dots, y'_{m,i+1}) \quad (20)$$

$$\bigwedge_{j=1}^m [a_j = 1 \wedge R_j(q_j, y_{j,1}, \dots, y_{j,i}) = 1] \vee [a_j = 0 \wedge R_j(q_j, y'_{j,2}, \dots, y'_{j,i+1}) = 0]. \quad (21)$$

In particular, (21) is our Σ_{i+1}^p formula for L . □

□

□

3.6.1 Exercises

(Recommended) Problem 49. Here, we prove Lemma 137. Let $i \geq 1$, and suppose that $\Sigma_i^p = \text{NP}^{\Sigma_{i-1}^p}$. We show that $\Sigma_{i+1}^p \subseteq \text{NP}^{\Sigma_i^p}$. Let $L \in \Sigma_{i+1}^p$, and let M be the Σ_i^p Turing Machine that accepts L . Suppose that $\omega \in L$. We have that:

$$\exists x_1, \forall x_2, \dots, Qx_{i+1} M(\omega; x_1, \dots, x_{i+1}) = 1,$$

where the x_i are all of length at most $\text{poly}(|\omega|)$. Define:

$$L' = \{\langle \omega, x_1 \rangle : \forall x_2, \dots, Qx_{i+1} M(\omega; x_1, \dots, x_{i+1}) = 1\}.$$

Do the following.

- (a) Show that $\langle \omega, x_1 \rangle \notin L'$ if and only if

$$\exists x_2, \dots, \overline{Q}x_{i+1} M(\omega; x_1, \dots, x_{i+1}) = 0.$$

- (b) Show that $\overline{L'} \in \Sigma_i^p$.

- (c) Design an $\text{NP}^{\Sigma_i^p}$ Turing Machine to accept L .

- (d) Conclude that $\Sigma_{i+1}^p \subseteq \text{NP}^{\Sigma_i^p}$.

(Recommended) Problem 50. We showed that for $i \geq 0$, $\Sigma_{i+1}^p = \text{NP}^{\Sigma_i^p}$. By definition, $\Pi_{i+1}^p = \text{co}\Sigma_{i+1}^p$. So $\Pi_{i+1}^p = \text{coNP}^{\Sigma_i^p}$. Show that: $\Pi_{i+1}^p = \text{coNP}^{\Pi_i^p}$.

3.7 Time Hierarchy Theorem

It is often easy to show that complexity classes \mathcal{C}_1 and \mathcal{C}_2 satisfy $\mathcal{C}_1 \subseteq \mathcal{C}_2$. Determining whether $\mathcal{C}_1 = \mathcal{C}_2$ is significantly harder. For instance, it is quite easy to show that $\text{P} \subseteq \text{NP}$. However, determining whether $\text{P} = \text{NP}$ is the central open problem in Theoretical Computer Science. In this section, we explore tools which allow us to separate certain complexity classes. We begin with the Time Hierarchy Theorem.

Definition 139. We say that a function $f : \mathbb{N} \rightarrow \mathbb{N}$ is *time-constructible* if $f(n) \geq n$ for all $n \in \mathbb{N}$. Similarly, we say that f is *space-constructible* if $f(n) \geq \log(n)$ for all $n \in \mathbb{N}$.

Remark 140. When dealing with time-complexity, we often assume that our runtime functions are time-constructible to ensure that the Turing Machines can read the entire inputs. We similarly assume that our space-complexity functions are space-constructible.

Before introducing the Time Hierarchy Theorem, we introduce the following fact.

Theorem 141. Let M be a k -tape Turing Machine. If M takes T steps to process a string ω , then a two-tape Turing Machine can simulate M in $O(T \log(T))$ steps.

Theorem 142 (Time Hierarchy Theorem). Suppose that $f, g : \mathbb{N} \rightarrow \mathbb{N}$ be time-constructible functions such that $f(n) \log(f(n)) \in o(g(n))$. Then $\text{DTIME}(f(n)) \subsetneq \text{DTIME}(g(n))$.

Proof. Clearly, $\text{DTIME}(f(n)) \subseteq \text{DTIME}(g(n))$. We use diagonalization to construct a language $L \in \text{DTIME}(g(n))$ such that $L \notin \text{DTIME}(f(n))$. We first recall that Turing Machines can be encoded as strings. Let $(w_n)_{n \in \mathbb{N}}$ be an enumeration of deterministic Turing Machine encodings, where each M_n is given by its corresponding encoding $\omega_n := \langle M_n \rangle$. We may also encode our Turing Machines in such a way as to allow for redundancy. Namely, for each Turing Machine M and each $k \in \mathbb{N}$, we may have the encoding $\langle M \rangle 1^k$. Using such an encoding scheme, any given Turing Machine will appear infinitely many times in our enumeration.

On input ω_n , we simulate M_n on input ω_n . We note that M_n may have more than two tapes. By Theorem 141, we may simulate M_n using only two tapes with a multiplicative overhead of $O(\log(n))$. We also note that there is a cost of converting the symbols used by M_n into a fixed, standard alphabet. So each step of M_n can be simulated in time $c \cdot T(n) \log(T(n))$, where $T(n)$ is the runtime complexity function for M_n . In particular, if $T(n) = f(n)$, then M_n can be simulated in time $O(f(n) \log(f(n)))$.

Now define:

$$L = \{\omega_n : M_n(\omega_n) \text{ halts after } g(|\omega_n|) \text{ steps and } M_n \text{ rejects } \omega_n\}.$$

We claim that $L \in \text{DTIME}(g(n))$ and $L \notin \text{DTIME}(f(n))$.

Lemma 143. $L \in \text{DTIME}(g(n))$.

Proof. On a given string ω , we may check whether ω is a valid Turing Machine encoding; and if so, simulate the corresponding Turing Machine M_ω for at most $g(|\omega|)$ steps. We accept ω if and only if M_ω halts and rejects ω . Otherwise, we reject ω . \square

Lemma 144. $L \notin \text{DTIME}(f(n))$.

Proof. Suppose to the contrary that $L \in \text{DTIME}(f(n))$. Then there exists a deterministic Turing Machine M such that $L(M) = L$ and M runs in time $O(f(n))$. Let ω be an encoding of M such that $f(|\omega|) \log(f(|\omega|)) < g(|\omega|)$. We analyze whether $\omega \in L$ to obtain our contradiction. We leave the details as an exercise. \square

\square

3.7.1 Exercises

(Recommended) Problem 51. Complete the proof of Lemma 144. Let ω be an encoding of M such that $f(|\omega|) \log(f(|\omega|)) < g(|\omega|)$ (such an encoding exists as $f(n) \log(f(n)) \in o(g(n))$), the $\text{DTIME}(f(n))$ Turing Machine that accepts L . We ask whether $\omega \in L$.

- (a) Suppose that $\omega \in L$. As M accepts L , we have that M accepts ω in $f(|\omega|)$ steps. Using the definition of L , deduce that $\omega \notin L$ to obtain a contradiction.

- (b) Suppose instead that $\omega \notin L$. Using a similar argument as part (a), deduce that $\omega \in L$ to obtain a contradiction.
- (c) Conclude that $L \notin \text{DTIME}(f(n))$.

(Recommended) Problem 52. Define:

$$\text{EXPTIME} := \bigcup_{c \in \mathbb{N}} \text{DTIME}(2^{n^c}).$$

Do the following.

- (a) Show that $\text{P} \subseteq \text{DTIME}(n^{\log(n)})$.
- (b) Fix $c \geq 1$. Show that $\text{DTIME}(n^{\log(n)}) \subsetneq \text{DTIME}(2^{n^c})$.
- (c) Deduce that $\text{P} \neq \text{EXPTIME}$.

Remark 145. We note that $\text{P} \subseteq \text{NP} \subseteq \text{PSPACE} \subseteq \text{EXPTIME}$. In light of Problem 52, we note that one of these containments is strict, though we don't know which one. It is believed that all of these containments are strict.

(Recommended) Problem 53. Let $\epsilon > 0$. Show that $\text{DTIME}(n^k) \subsetneq \text{DTIME}(n^{k+\epsilon})$. That is, there are problems in P that require arbitrarily large exponents to solve.

(Recommended) Problem 54. In this problem, we prove the Space Hierarchy Theorem, which states that if $f, g : \mathbb{N} \rightarrow \mathbb{N}$ are space-constructible functions satisfying $f(n) \in o(g(n))$, then $\text{DSPACE}(f(n)) \subsetneq \text{DSPACE}(g(n))$. As in the proof of the Time Hierarchy Theorem, let $(w_n)_{n \in \mathbb{N}}$ be an enumeration of deterministic Turing Machine encodings, where each M_n is given by its corresponding encoding $\omega_n := \langle M_n \rangle$. Note that Turing Machine encodings are not unique. It is in fact the case that any given Turing Machine will appear infinitely many times in our enumeration. Define:

$$L = \{\omega_n : M_n(\omega_n) \text{ halts and rejects } \omega_n \text{ using space } g(|\omega_n|) \text{ and taking at most } 2^{g(n)} \text{ steps.}\}.$$

Do the following.

- (a) Why do we clock the execution of $M_n(\omega_n)$ to $2^{g(n)}$ steps?
- (b) Show that $L \in \text{DSPACE}(g(n))$.
- (c) We now show that $L \notin \text{DSPACE}(f(n))$. Suppose to the contrary that there exists a $\text{DSPACE}(f(n))$ Turing Machine M that decides L . Let ω be an encoding of M such that $f(|\omega|) < g(|\omega|)$ (such an encoding exists as $f(n) \in o(g(n))$).
 - (i) Suppose that $\omega \in L$. As M accepts L , we have that M accepts ω using at most $g(|\omega|)$ space and taking at most $2^{g(|\omega|)}$ steps. Using the definition of L , deduce that $\omega \notin L$ to obtain a contradiction.
 - (ii) Suppose that $\omega \notin L$. Using a similar argument as in part (a), deduce that $\omega \in L$ to obtain a contradiction.
 - (iii) Conclude that $L \notin \text{DSPACE}(f(n))$.

(Recommended) Problem 55. Using the Space Hierarchy Theorem (Problem 54), show that $\text{L} \subsetneq \text{DSPACE}(n)$. Conclude that $\text{L} \subsetneq \text{PSPACE}$.

Remark 146. We have that:

$$\text{L} \subseteq \text{NL} \subseteq \text{AC}^1 \subseteq \text{NC}^2 \subseteq \dots \subseteq \text{NC} \subseteq \text{P} \subseteq \text{NP} \subseteq \text{PSPACE}.$$

In light of Problem 55, we have that one of these containments is strict. It is widely believed that all of these containments are strict.

4 Interactive Proofs

An interactive proof system is an abstract computational machine that permits interaction between two parties: a prover and a verifier. For a given language L , the prover seeks to convince the verifier whether the string $\omega \in L$. In order to convince the verifier, the prover offers a *certificate* or *proof*. The verifier then examines the proof and makes a decision as to whether the proof is valid. The initial proof may not be convincing, and so the verifier may have questions for the prover. The prover is then provided a chance to respond. This could occur for multiple iterations or rounds. This is much like the process of engaging in mathematical research or in submitting a paper to a journal.

In practice as well, the verifier need not be perfect. That is, the verifier could erroneously reject a correct proof or accept an incorrect proof. When the verifier behaves deterministically and without error, we capture precisely NP. We may generalize this process in two ways: by allowing the verifier to make mistakes, or by introducing randomness into the process. It turns out that the combination of interaction and error is quite powerful, characterizing PSPACE. Our discussions of randomness will take us into a different direction: namely Arthur-Merlin protocols.

4.1 Preliminaries

Intuitively, we want three conditions to hold in a theorem-proving procedure.

- It is possible to prove a true theorem.
- It is impossible to prove a false theorem.
- It does not matter how long it takes for the prover to produce a proof, but it must be easy for the verifier to check that the proof is correct.

In a static proof system, a proof is submitted. The verifier then examines the proof and makes a decision as to whether the proof is valid. However, another method of convincing a verifier is through interaction. That is, the verifier may have questions for the prover, and the prover has a chance to respond. This could occur for multiple iterations or rounds. We formalize the notion of interaction as follows. Note that either the prover or the verifier may go first. If the verifier goes first, we may think of this intuitively as the verifier posing a problem for the prover to solve.

Definition 147 (Interaction of Deterministic Functions). Let $f, g : \{0, 1\}^* \rightarrow \{0, 1\}^*$ be functions, and let $k \in \mathbb{N}$. A k -round interaction of f and g on input $x \in \{0, 1\}^*$ is the following sequence of strings a_1, \dots, a_k defined as follows:

$$\begin{aligned} a_1 &= f(\langle x \rangle) \\ a_2 &= g(\langle x, a_1 \rangle) \\ a_3 &= f(\langle x, a_1, a_2 \rangle) \\ a_4 &= g(\langle x, a_1, a_2, a_3 \rangle) \\ &\dots \\ a_{2i+1} &= f(\langle x, a_1, \dots, a_{2i} \rangle) \\ a_{2i+2} &= g(\langle x, a_1, \dots, a_{2i+1} \rangle). \end{aligned}$$

The output of f, g are denoted $\text{out}_f(\langle f, g \rangle), \text{out}_g(\langle f, g \rangle)$.

Definition 148 (Deterministic Interactive Proofs). We say that a language L has a k -round *deterministic interactive proof system* if there exists a deterministic TM V (which we call our verifier) that, on input $\langle x, a_1, a_2, a_3, \dots, a_k \rangle$, runs in polynomial time in $|x|$ and satisfies the following:

- **Completeness:** If $x \in L$, then there exists a prover $P : \{0, 1\}^* \rightarrow \{0, 1\}^*$ such that $\text{out}_V(\langle P, V \rangle) = 1$.
- **Soundness:** If $x \notin L$, then for every prover $P : \{0, 1\}^* \rightarrow \{0, 1\}^*$, $\text{out}_V(\langle P, V \rangle) = 0$.

We note that as V runs in polynomial time on $|x|$, we may assume without loss of generality that for each i , $|a_i|$ is polynomial in $|x|$.

Remark 149. Here, $\langle a_1, \dots, a_k \rangle$ is the transcript arising from the interaction of V and the prover P providing the correct proof. Note that as V is deterministic, there is a single accepting computation for any given input string. It follows that there is a unique, correct proof for any string x in the given language L .

Definition 150. The class **dIP** is the set of languages L , where there exists a $k(n)$ -round deterministic interactive proof system, where $k(n)$ is a fixed polynomial depending on L . Here, the input is $n = |x|$.

The first result we wish to prove is that **dIP** = **NP**. Before doing so, we first recall the verifier definition of **NP**.

Definition 151 (NP). We say that the language $L \in \text{NP}$ if there exists a polynomial time deterministic TM V such that: $x \in L$ if and only if there exists a certificate C (where $|C|$ is polynomial in $|x|$) such that $V(x, C) = 1$. Here, V is our verifier.

Remark 152. In light of the verifier definition of **NP**, we see that every language L has a 1-step deterministic interactive proof system. Here, we view $a_1 = C$.

Theorem 153. **dIP** = **NP**.

Proof. We first show that **NP** \subseteq **dIP**. Let $L \in \text{NP}$, and let V be the verifier for L as in the definition of **NP**. Observe that V also satisfies the definition of a 1-round deterministic interactive proof system for L . So $L \in \text{dIP}$.

We next show that **dIP** \subseteq **NP**. Let $L \in \text{dIP}$, and let V be the verifier. We show that $x \in L$ if and only if the corresponding transcript $\langle a_1, a_2, \dots, a_k \rangle$ causing V to accept serves as the certificate, as in the definition of **NP**. We leave the remainder of the proof as an exercise for the reader (see Exercise 153). The key idea is that the transcript must be polynomial in length, as the interactive proof protocol has only polynomially-many rounds (see the definition of **dIP**; Definition 150). \square

4.1.1 Exercises

(Recommended) Problem 56. Complete the proof of Theorem 153 by showing that **dIP** \subseteq **NP**.

4.2 Complexity Class IP

In the previous section, we examined deterministic interactive proof systems, in which the verifier behaves perfectly. In practice, the verifier is often imperfect. That is, the verifier could erroneously reject a correct proof or accept an incorrect proof. It turns out that this combination of interaction and randomness is quite powerful, capturing PSPACE instead of just NP. There are two types of errors.

- **Completeness:** The verifier could reject a valid proof.
- **Soundness:** The verifier could accept an incorrect proof.

This leads us to the definition of the complexity class IP.

Definition 154. A language $L \in \text{IP}$ if there exist interactive algorithms $\langle P, V \rangle$ with V running in probabilistic polynomial time (in the length of the common input x), such that:

- **Completeness:** If $x \in L$, then $\Pr[\text{out}_V(\langle P, V \rangle)(x) = 1] \geq \frac{2}{3}$.
- **Soundness:** If $x \notin L$, then for any prover P^* , $\Pr[\text{out}_V(\langle P^*, V \rangle)(x) = 1] \leq \frac{1}{3}$.

Remark 155. Note that the provers P, P^* can be computationally unbounded. It doesn't matter how long it takes the prover to produce a proof, only that it is correct and easy to verify.

We now look at examples of languages that belong to IP.

Example 156. Recall that the Graph Isomorphism problem (GI) takes as input two graphs, G_0 and G_1 , and asks whether $G_0 \cong G_1$. We show that Graph Isomorphism $\in \text{IP}$. The interactive proof protocol has k iterations (we will discuss how to determine k later). At each iteration, the following occurs.

- Verifier generates a random bit $b \in \{0, 1\}$ and a permutation π of the graph G_b . Next, Verifier sends $H := \pi(G_b)$ to Prover.
- Prover responds with a bit $b' \in \{0, 1\}$ and a permutation π' such that $\pi'(H) = G_{b'}$. That is, Prover asserts that $\pi' = \pi^{-1}$.
- Verifier checks whether π' is indeed an isomorphism.

Verifier accepts if Prover was able to send a correct permutation and bit at each round. We now analyze the Completeness and Soundness properties of this protocol.

- **Completeness:** If $G_0 \cong G_1$, then any bit b' and any permutation π' that Prover sends will be an isomorphism of $G_{b'}$. So valid proofs are accepted with probability 1.
- **Soundness:** Suppose that $G_0 \not\cong G_1$. At a given round, if Prover's proof (that is, the pair (b', π') , where π' is claimed to be isomorphism of H and $G_{b'}$) is invalid, then Verifier catches this with probability $1/2$. So with k rounds, the probability that Prover will have an invalid proof accepted is at most 2^{-k} . In order to ensure that this probability is at most $1/3$, we require $k = 2$ rounds.

4.2.1 Exercises

(Recommended) Problem 57. Modify Example 156 to show that Graph Non-Isomorphism $\in \text{IP}$.

(Recommended) Problem 58. Let $m \in \mathbb{Z}^+$. We say that y is a *quadratic residue* modulo m if there exists an x such that $y \equiv x^2 \pmod{m}$. Show that deciding whether y is **not** a quadratic residue modulo m is in IP. Precisely, show that the following language is in IP:

$$\text{QNR} = \{(y, m) : y \text{ is not a quadratic residue modulo } m.\}$$

Carefully determine the number of rounds needed in the protocol.

4.3 IP = PSPACE

It turns out that $\text{IP} = \text{PSPACE}$.

Theorem 157. $\text{IP} = \text{PSPACE}$.

We first show that $\text{IP} \subseteq \text{PSPACE}$.

Proposition 158. $\text{IP} \subseteq \text{PSPACE}$.

Proof. Let $L \in \text{IP}$. Let V be the probabilistic polynomial time verifier for L . Our goal is to design a PSPACE algorithm to compute:

$$z := \max_P \Pr[\text{out}_V(\langle P, V \rangle)(w) = 1],$$

where $w \in \{0, 1\}^*$ is arbitrary and we are taking a maximum over all provers P . From the definition of IP , we have that:

- If $z \geq \frac{2}{3}$, then $w \in L$; and
- If $z \leq \frac{1}{3}$, then $w \notin L$.

As V is a verifier for L , the case that $\frac{1}{3} < z < \frac{2}{3}$ will never occur. We now provide a PSPACE algorithm to compute z . We note that V runs in polynomial time, for a given polynomial $p(n)$ (where $n = |w|$). So any response provided by the prover must be of length at most $p(n)$. As V is probabilistic, it picks at most one random number at each step (where each random number selected represents a corresponding branching step). As V chooses at most $p(n)$ random numbers, we only need a polynomial amount of space to represent the choices. So the accepting computations can be enumerated using a polynomial amount of space. Now z is the total number of accepting branches, divided by the total number of possible branches. Thus, $L \in \text{PSPACE}$, and we conclude that $\text{IP} \subseteq \text{PSPACE}$. \square

Proposition 159. $\text{PSPACE} \subseteq \text{IP}$.

Proof. Recall that TQBF is PSPACE-complete, even when the formulas are restricted to be in 3-CNF. We show that $3\text{CNF-TQBF} \in \text{IP}$. Let ψ be an instance of 3CNF-TQBF . We first encode φ into a polynomial, using a process referred to as *arithmetization*. We then design an interactive proof protocol to analyze the polynomials.

Step 1: We construct a polynomial φ corresponding to ψ . We begin by arithmetizing the non-quantified Boolean expression, which provides a polynomial of degree $3m$, where m is the number of clauses. The arithmetization is as follows:

$$\begin{aligned} x_i &\mapsto z_i \\ \bar{x}_i &\mapsto (1 - z_i) \\ x \vee y \vee z &= 1 - (1 - x)(1 - y)(1 - z) \\ x \wedge y &= x \cdot y. \end{aligned}$$

Denote Φ to be the arithmetization of the Boolean formula (without the quantifiers). We now describe how to arithmetize the quantifiers.

- The existential quantifier $\exists x_i \phi(x_1, \dots, x_n)$ is arithmetized as follows:

$$\prod_{x_i \in \{0, 1\}} \Phi(x_1, \dots, x_n) = 1 - \left[\left(1 - \Phi(x_1, \dots, x_{i-1}, 0, x_{i+1}, \dots, x_n) \right) \cdot \left(1 - \Phi(x_1, \dots, x_{i-1}, 1, x_{i+1}, \dots, x_n) \right) \right]$$

Note that in order for ψ to be true, ϕ must be true for some value of x_i . If $x_i = 0$ accomplishes this, then the first term is 0, making the expression above 1. The same result holds for the second term if instead $x_i = 1$ accomplishes this.

- Now the universal quantifier $\forall x_i \phi(x_1, \dots, x_n)$ is arithmetized as follows:

$$\prod_{x_i \in \{0, 1\}} \Phi(x_1, \dots, x_n) = \Phi(x_1, \dots, x_{i-1}, 0, x_{i+1}, \dots, x_n) \cdot \Phi(x_1, \dots, x_{i-1}, 1, x_{i+1}, \dots, x_n)$$

Note that this ensures that the condition holds for all values of x_i .

Observe that arithmetizing a universal or existential quantifier doubles the degree of the final polynomial. So in the worst case, the final polynomial Φ has degree $2^n \cdot 3m$, where we have n quantifiers and m clauses. As the verifier V must run in polynomial time, it cannot parse a polynomial with exponentially many terms, making the degree of the polynomial too large for the verifier to handle.

Note that we are only interested in the values that the x_i 's take on in $\{0, 1\}$, we have that $x_i^k = x_i$ for every $k > 0$. We use this observation to reduce the degree of our polynomial. Let R_{x_i} denote the degree reduction operator for the variable x_i . Precisely:

$$R_{x_i}(\Phi(x_1, \dots, x_n)) = x_i \Phi(x_1, \dots, x_{i-1}, 1, x_{i+1}, \dots, x_n) + (1 - x_i) \Phi(x_1, \dots, x_{i-1}, 0, x_{i+1}, \dots, x_n).$$

So $R_{x_i}(\Phi)$ agrees with Φ on all values of $x_i \in \{0, 1\}$, and each x_i term in $R_{x_i}(\Phi)$ has degree at most 1. Thus, our polynomial can be written as:

$$\varphi = \prod_{x_1} R_{x_1} \prod_{x_2} R_{x_2} \prod_{x_3} \dots \prod_{x_n} R_{x_n} \Phi(x_1, \dots, x_n).$$

We have $\mathcal{O}(n^2)$ invocations of a reduction operator R_{x_i} , and each R_{x_i} can be computed in polynomial time. So the degree reduction is polynomial time computable.

Observe that $\varphi = 1$ if and only if ψ is true.

Step 2: The goal of the Prover is to convince the Verifier that $\varphi = 1$. As $\varphi = 0$ or $\varphi = 1$, we may (WLOG) evaluate φ over \mathbb{F}_q for any prime q (though the soundness probability depends inversely on q ; we will discuss this later). So it suffices for the Prover to convince the Verifier that:

$$\varphi \equiv 1 \pmod{q}.$$

Abstractly, we may write:

$$\varphi = \mathcal{O}_1 \mathcal{O}_2 \dots \mathcal{O}_\ell \Phi(x_1, \dots, x_n) \pmod{q},$$

where each \mathcal{O}_i is either a \prod_{x_i} , \prod_{x_i} , or R_{x_i} . Note that:

$$\ell = \sum_{i=1}^n (i + 1).$$

The idea of the proof protocol is to strip off one of the \mathcal{O}_i terms at each round. We proceed as follows.

- The verifier sets $v_0 = 1$ and $\Phi_0 = \Phi$. So the prover wants to convince the verifier that the given quantified formula is true.
- We note that \mathcal{O}_1 is either a \prod or \prod . The prover sends a degree-1 polynomial $\hat{p}(x_1)$. The verifier checks that $v_0 = 1 = \prod \hat{p}(x_1)$ or $v_0 = 1 = \prod \hat{p}(x_1)$ respectively. Should this condition fail to hold, the verifier rejects. Otherwise, the verifier selects a random $r_1 \in \mathbb{F}_q$ and sets $v_1 := \hat{p}(r_1)$, and:

$$\Phi_1 = \Phi(r_1, x_2, \dots, x_n).$$

We then enter the next round.

More Generally, the protocol is defined as follows:

- The verifier sets $v_0 = 1$ and $\Phi_0 = \Phi$. So the prover wants to convince the verifier that the given quantified formula is true.
- At round k , the verifier has some value v_k , as well as some polynomial Φ_k . The prover's goal is to convince the verifier that:

$$v_k = \mathcal{O}_{k+1} \dots \mathcal{O}_\ell \Phi(x_1, \dots, x_n) \pmod{q}. \tag{22}$$

The verifier then computes v_{k+1} and Φ_{k+1} , and we proceed to the next round. We describe the cases for each round.

- **Case 1:** $\mathcal{O}_{k+1} = \prod_{x_i}$. Here, prover wishes to convince the verifier that:

$$v_k = [\mathcal{O}_{k+1} \cdots \mathcal{O}_\ell \Phi(x_1, \dots, x_n)] \Big|_{r_1, \dots, r_{i-1}} \pmod{q} \quad (23)$$

Note that the expression is evaluated symbolically, and then the r_j 's are plugged into the expression. The protocol proceeds as follows:

- The prover sends some degree-1 polynomial $\hat{p}(x_i)$.
- The verifier checks that $v_k = \hat{p}(0) \cdot \hat{p}(1)$. If this fails to hold, the verifier rejects. Otherwise, the verifier selects a random $r_i \in \mathbb{F}_q$ and sets $v_{k+1} = \mathcal{O}_{k+2} \cdots \mathcal{O}_\ell \Phi(r_1, \dots, r_i, x_{i+1}, \dots, x_n)$. Now $\Phi_k := \Phi(r_1, \dots, r_i, x_{i+1}, \dots, x_n)$.

We now verify completeness and soundness.

- **Completeness:** Suppose equation (22) holds. Then the prover can send

$$P(x_i) := \mathcal{O}_{k+1} \cdots \mathcal{O}_\ell \Phi(x_1, \dots, x_n).$$

The verifier will not reject, regardless of the choice of r_i .

- **Soundness:** If equation (22) does not hold, then the prover must send some polynomial $\hat{P}(x_i) \neq P(x_i)$ (where $P(x_i)$ was defined in Completeness). Otherwise, the verifier rejects immediately. As \hat{P} and P are degree-one polynomials, they can agree on at most one point $r \in \mathbb{F}_q$. So (23) will not hold, except with probability $1/q$.

- **Case 2:** $\mathcal{O}_{k+1} = \prod_{x_i}$. This case is analogous to the previous case, except that we are checking:

$$v_k = \prod \hat{p}(x_i) = 1 - (1 - p(0)) \cdot (1 - p(1)) \pmod{q}.$$

Lemma 160. Suppose equation (22) holds. Then the verifier will accept Prover's proof at round k with probability 1.

Proof. Exercise. □

Lemma 161. Suppose equation (22) does not hold. Then the verifier will accept Prover's proof at round k with probability at most $1/q$.

Proof. Exercise. □

- **Case 3:** $\mathcal{O}_{k+1} = R_{x_i}$. Note that as we are considering R_{x_i} , this operator must appear after some \prod_{x_j} or \prod_{x_j} , where $j \geq i$. Thus, the prover wants to convince the verifier that:

$$v_k = R_{x_i} \cdots \mathcal{O}_\ell \Phi(r_1, \dots, r_j, x_{j+1}, \dots, x_n) \pmod{q}. \quad (24)$$

The proof protocol works as follows.

- The prover sends a polynomial $\hat{p}(x_i)$. The degree of this polynomial is at most $3m$ if the R_{x_i} is one of the inner-most reduction operators (as Φ is a polynomial of degree at most $3m$). Otherwise, \hat{p} has degree at most 2.
- The verifier checks that $v_k = (R_{x_i} \hat{p})[r_i] \pmod{q}$ (**Note:** As $j \geq i$, r_i was selected on a previous round). If this condition does not hold, the verifier rejects. Otherwise, the verifier chooses a **new** random $r'_i \in \mathbb{F}_q$ and sets:

$$v_{k+1} := \mathcal{O}_{k+2} \cdots \mathcal{O}_\ell \Phi(r_1, \dots, r_{i-1}, r'_i, r_{i+1}, \dots, r_j, x_{j+1}, \dots, x_n) \pmod{q}. \quad (25)$$

Lemma 162. Suppose equation (24) holds. Then the verifier will accept Prover's proof at round k with probability 1.

Proof. Exercise. □

Lemma 163. Suppose that equation (24) does not hold.

- (a) If R_{x_i} is one of the n inner-most reduction operators, then Verifier accepts Prover's proof with probability at most $3m/q$.
- (b) If R_{x_i} is not one of the n inner-most reduction operators, then Verifier accepts Prover's proof with probability at most $2/q$.

Proof. Exercise. □

We now examine the Completeness and Soundness bounds for the entire protocol.

- **Completeness of Protocol:** Each case has completeness probability 1. So a correct proof is accepted with probability 1.
- **Soundness for Protocol:** There are n operators of the form Π, II , each of which contributes error $1/q$. Now the interior n operators of the form R_{x_i} each contribute error at most $3m/q$, for a total error bound of $3mn/q$. The remaining R_{x_i} operators each contribute error $2/q$. There are $\sum_{i=1}^{n-1} i$ such operators. So the total error bound is:

$$\frac{n}{q} + \frac{3mn}{q} + \frac{2}{q} \sum_{i=1}^{n-1} i = \frac{3mn + n^2}{q}$$

We may select a q of polynomial length to obtain a soundness error of at most $1/3$, as desired. □

□

4.3.1 Exercises

(Recommended) **Problem 59.** Prove Lemma 160.

(Recommended) **Problem 60.** Prove Lemma 161.

(Recommended) **Problem 61.** Prove Lemma 162.

(Recommended) **Problem 62.** Prove Lemma 163.

4.4 Public Coins and Arthur-Merlin Protocols

The interactive proof protocol for Graph Isomorphism (Example 156) relied heavily on the fact that the selections made by the Verifier were private; that is, the information was unavailable to the prover. It is natural to ask as to the power of interactive proof protocols where the Verifier's choices are made public. This motivates the notion of Arthur-Merlin protocols. Here, Merlin takes the role of the all-powerful Prover, while Arthur takes the role of the Verifier. We begin with the definition of the classes $\text{AM}[k]$ and $\text{MA}[k]$.

Definition 164. We say that $L \in \text{AM}[k]$ if L has a k -round Arthur-Merlin protocol with Arthur making the first move. Here, Arthur selects his strings arbitrarily, and Merlin must be able to respond with an appropriate answer. We denote Arthur selecting a string as Ax_i , and we denote Merlin selecting a string as Mx_i . Let $x \in \Sigma^*$. After the k rounds, Arthur decides whether to accept the string x . We have the following conditions.

- **Completeness:** If $x \in L$, then:

$$Aa_1, Ma_2 Aa_3 Ma_4 \dots, Qa_k \Pr[\text{Arthur}(\langle x, a_1, \dots, a_k \rangle) = 1] = 1.$$

- **Soundness:** If $x \notin L$, then:

$$Aa_1, Ma_2 Aa_3 Ma_4 \dots, Qa_k \Pr[\text{Arthur}(\langle x, a_1, \dots, a_k \rangle) = 1] < 1/4.$$

We similarly define $\text{MA}[k]$.

Definition 165. We say that $L \in \text{MA}[k]$ if L has a k -round Arthur-Merlin protocol with Merlin making the first move. Here, Arthur selects his strings arbitrarily, and Merlin must be able to respond with an appropriate answer. We denote Arthur selecting a string as Ax_i , and we denote Merlin selecting a string as Mx_i . Let $x \in \Sigma^*$. After the k rounds, Arthur decides whether to accept the string x . We have the following conditions.

- **Completeness:** If $x \in L$, then:

$$Ma_1, Aa_2 Ma_3 Aa_4 \dots, Qa_k \Pr[\text{Arthur}(\langle x, a_1, \dots, a_k \rangle) = 1] = 1.$$

- **Soundness:** If $x \notin L$, then:

$$Ma_1, Aa_2 Ma_3 Aa_4 \dots, Qa_k \Pr[\text{Arthur}(\langle x, a_1, \dots, a_k \rangle) = 1] < 1/4.$$

Remark 166. We note that the nested quantifier formulations of $\text{AM}[k]$ and $\text{MA}[k]$ resemble that of the Polynomial-Time Hierarchy. In particular, the Arthur quantifier is analogous to the universal quantifier, and the Merlin quantifier is analogous to the existential quantifier. To see this, regardless of Arthur's random selections, Merlin need only find one appropriate response to convince Arthur to accept. So $\text{AM}[k]$ is analogous to Π_k^p , and $\text{MA}[k]$. This raises a couple questions.

- How are $\text{AM}[k]$ and $\text{MA}[k]$ related, if at all? It is clear that for all $k \in \mathbb{N}$, $\text{AM}[k] \subseteq \text{AM}[k+1]$ and $\text{MA}[k] \subseteq \text{MA}[k+1]$. In particular, do $\text{AM}[k]$ and $\text{MA}[k]$ play analogous roles as Π_k^p and Σ_k^p in the Polynomial-Time Hierarchy?
- Do either the AM-hierarchy or MA hierarchy collapse? That is, does there exist a k such that $\text{AM}[j] \subseteq \text{AM}[k]$ for all $j \geq k$? Similarly, does there exist a k such that $\text{MA}[j] \subseteq \text{MA}[k]$ for all $j \geq k$?
- Are $\text{AM}[k]$ and $\text{MA}[k]$ related to Π_k^p and Σ_k^p ?

Remark 167. We also remark on some terminology. Namely, we refer to the Verifier's selections as *coins*. Instances where Verifier's coins are not known to Prover are referred to as *private coins*, whereas instances where the Verifier's selections are known to Prover are referred to as *public coins*.

We show that $\text{MA}[k] \subseteq \text{AM}[k]$. We begin by showing that $\text{MA}[2] \subseteq \text{AM}[2]$. Our proof that $\text{MA}[2] \subseteq \text{AM}[2]$ is adopted from [BPG04].

Remark 168. Before proving that $\text{MA}[2] \subseteq \text{AM}[2]$, we pause to discuss intuitively why it should be the case that an AM protocol with k rounds is at least as powerful as an MA protocol with k rounds. In the first round of an MA protocol, Merlin makes a move. We note that Merlin's selection is existentially quantified. If Merlin knows Arthur's move, then Merlin can more strategically select his response. This is precisely the setting of an AM protocol.

Theorem 169. $\text{MA}[2] \subseteq \text{AM}[2]$.

Proof. Let $L \in \text{MA}[2]$. Fix a string $x \in \Sigma^*$. Suppose that Merlin sends the message a_1 at the first round. Arthur responds with a random coin a_2 . We note that if $x \in L$, then Arthur accepts $\langle x, a_1, a_2 \rangle$ with probability 1; and if $x \notin L$, then Arthur accepts $\langle x, a_1, a_2 \rangle$ with probability strictly less than $1/4$.

We define a two-round AM protocol as follows. Let V be the protocol employed by Arthur in the $\text{MA}[2]$ protocol for L . At the initial round, Arthur selects m random coins r_1, \dots, r_m , where $m = \text{poly}(|x|)$ and each $|r_i| = \text{poly}(|x|)$. Arthur sends $\langle r_1, \dots, r_m \rangle$. Merlin then responds with the string y . Arthur accepts if $\langle x, \langle r_1, \dots, r_m \rangle, y \rangle$ precisely if:

$$\text{Majority}_{i \in [m]} V(x, y, r_i) = 1. \quad (26)$$

We now analyze the Completeness and Soundness of the $\text{AM}[2]$ protocol.

- **Completeness:** We note that Merlin may send the string $y = a_1$ as in the $\text{MA}[2]$ protocol. We note that if $x \in L$, $V(x, a_1, r) = 1$ for all r . So in the AM protocol defined above, we have that if $x \in L$, then Arthur accepts with probability 1.
- **Soundness:** Analyzing the Soundness of this protocol is more involved. Our goal is to show that (26) holds with probability strictly less than $1/2$. Denote:

$$H(x, y, r) := \Pr[V(x, y, r) = 1].$$

We note that:

$$\begin{aligned} \Pr \left[\bigwedge_{i=1}^m V(x, y, r_i) = 1 \right] &= \prod_{i=1}^m \Pr[V(x, y, r) = 1] \\ &= \prod_{i=1}^m H(x, y, r_i). \end{aligned}$$

Denote R to be the finite set of possible random choices for Arthur, and let Y denote the finite set of possible messages for Merlin. Fix $\bar{y} \in Y$. We bound the probability that (26) holds, by averaging over our choices of r_1, \dots, r_m . We note that:

$$\Pr[\text{Majority}_{i \in [m]} V(x, \bar{y}, r_i) = 1] \leq \frac{1}{|R|^m} \cdot \left(\sum_{\vec{r} \in R^m} \sum_{I \in \binom{[m]}{\lceil m/2 \rceil}} \prod_{i \in I} H(x, \bar{y}, r_i) \right) \quad (27)$$

$$= \frac{1}{|R|^m} \sum_{I \in \binom{[m]}{\lceil m/2 \rceil}} \sum_{\vec{r} \in R^m} \prod_{i \in I} H(x, \bar{y}, r_i) \quad (28)$$

We note that for a fixed I , only the terms $H(x, \bar{y}, r_i)$ where $i \in I$ contribute to the probability. So we may take an average over elements of $R^{\lceil m/2 \rceil}$ as opposed to R^m . It follows that the expression (28) is at most:

$$\begin{aligned} &\frac{1}{|R|^{\lceil m/2 \rceil}} \sum_{I \in \binom{[m]}{\lceil m/2 \rceil}} \sum_{\vec{r} \in R^{\lceil m/2 \rceil}} \prod_{i \in I} H(x, \bar{y}, r_i) \\ &= \sum_{I \in \binom{[m]}{\lceil m/2 \rceil}} \prod_{i \in I} \left(\sum_{r_i \in R} \frac{H(x, \bar{y}, r_i)}{|R|} \right) \\ &= \sum_{I \in \binom{[m]}{\lceil m/2 \rceil}} (Ar H(x, \bar{y}, r))^{\lceil m/2 \rceil}. \end{aligned}$$

Note that the expression $Ar H(x, \bar{y}, r)$ denotes the probability that Arthur accepts x when selecting the random coin r . Now let y^* be Merlin's best possible choice. We have the following, which corresponds to an MA[2] protocol, as Merlin chose y^* first.

$$\Pr[\text{Majority}_{i \in [m]} V(x, y^*, r_i) = 1] \leq \sum_{I \in \binom{[m]}{\lceil m/2 \rceil}} (Ar, \Pr[V(x, y^*, r) = 1])^{\lceil m/2 \rceil} \quad (29)$$

$$\leq 2^m (Ar, \Pr[V(x, y^*, r) = 1])^{\lceil m/2 \rceil} \quad (30)$$

$$= 2^m (My^*, Ar, \Pr[V(x, y^*, r) = 1])^{\lceil m/2 \rceil} \quad (31)$$

$$\leq 2^m \sum_{y \in Y} (Ar, \Pr[V(x, y^*, r) = 1])^{\lceil m/2 \rceil} \quad (32)$$

$$\leq 2^m |Y| (Ar, \Pr[V(x, y^*, r) = 1])^{\lceil m/2 \rceil} \quad (33)$$

Here, (32) follows by simulating V on all possible choices for Merlin. To obtain (33), we select an optimal y^* ; the $|Y|$ term follows from the number of ways to select a choice for Merlin. Denote $\Psi(x)$ to be the probability that an MA[2] protocol accepts x . We note that:

$$\Psi(x) := Ar, \Pr[V(x, y^*, r) = 1].$$

So in particular, (33) can be written more succinctly as:

$$2^m \cdot |Y| \cdot (\Psi(x))^{\lceil m/2 \rceil}. \quad (34)$$

Our goal now is to select m so that (34) is strictly less than $1/4$. So the probability of our AM[2] protocol accepting $x \notin L$ is at most:

$$2^m \cdot |Y| \cdot (\Psi(x))^{\lceil m/2 \rceil} = 2^m \cdot 2^{\log_2(\Psi(x))m/2} \cdot |Y|.$$

As $\Psi(x) < 1/4$ (this follows from the assumption that $x \notin L$, for the purposes of analyzing Soundness) for our analysis of, we have that $2^m \cdot 2^{\log_2(\Psi(x))m/2} = 2^{-\epsilon m}$ for some $\epsilon > 0$. Let $m := (4/\epsilon) \cdot \log_2(|Y|) + 4/\epsilon$. So we have that:

$$\begin{aligned} 2^{-\epsilon m} \cdot |Y| &= 2^{-\log_2(|Y|) - 4/\epsilon} \cdot |Y| \\ &= 2^{-4} \\ &= 1/16 \\ &< 1/4. \end{aligned}$$

So if $x \notin L$, our AM[2] protocol accepts x with probability at most $1/4$. □

Remark 170. Applying Theorem 169 and induction, we obtain the following corollary.

Corollary 171. For each constant $k \geq 2$, we have that $\text{MA}[k] \subseteq \text{AM}[k]$.

We also obtain as a corollary to Theorem 169 that $\text{AM}[k] \subseteq \text{AM}[2]$ for any $k > 2$. The idea is that we can use Theorem 169 to convert two rounds, starting with Merlin and ending with Arthur, into two rounds, starting with Arthur and ending with Merlin. Effectively, we go from three rounds with Arthur, Merlin, and then Arthur, to three rounds with Arthur, Arthur, and Merlin. The two Arthur rounds can effectively be combined into a single round.

Corollary 172. For each $k > 2$, we have that $\text{AM}[k] \subseteq \text{AM}[2]$.

Remark 173. In light of Corollary 172, we may unambiguously define $\text{AM} := \text{AM}[2] = \text{AM}[3]$.

Remark 174. It remains open whether $\text{MA}[2] = \text{AM}$.

We also note that constant round interactive proof protocols with private coins are no more powerful than if we allow for public coins. We direct the reader to the corresponding paper of Goldwasser and Sipser [GS86] for the proof.

Theorem 175 (Goldwasser-Sipser [GS86]). Let $p(n)$ be a polynomial. We have that $\text{IP}[p(n)] \subseteq \text{AM}[p(n) + 2]$.

Corollary 176. $\text{IP}[O(1)] = \text{AM}$.

4.4.1 Exercises

(Recommended) Problem 63. Let $\delta \in (0, 1/4)$ be a fixed constant. Strengthen the Soundness bound in Theorem 169 to show that $MA[2] \subseteq AM[2]$, when the soundness bound for the $AM[2]$ protocol is stipulated to be δ rather than $1/4$.

(Recommended) Problem 64. Prove Corollary 171: for each constant $k \geq 2$, we have that $MA[k] \subseteq AM[k]$.

(Recommended) Problem 65. Prove Corollary 172: for each $k > 2$, we have that $AM[k] \subseteq AM[2]$.

(Recommended) Problem 66. In Problem 57, we established that $\text{Graph Non-Isomorphism} \in IP[2]$. Deduce that $\text{Graph Non-Isomorphism} \in AM$.

4.5 Arthur-Merlin Protocols and the Polynomial-Time Hierarchy

In the previous section (see Remark 166), we established an analogue between Arthur-Merlin protocols and the quantifier formulation of the Polynomial-Time Hierarchy. Namely, Arthur selecting random bits correspond to universal quantifiers, and Merlin's selections correspond to existential quantifiers. As $\text{AM} := \text{AM}[2] = \text{AM}[k]$ for each $k \geq 2$, it is natural to ask where (if at all) AM is contained. As $\text{MA}[2] \subseteq \text{AM}$, understanding where AM lies within PH provides insight as to where $\text{MA}[2]$.

We begin by showing that $\text{AM} \subseteq \Pi_2^p$.

Theorem 177. $\text{AM} \subseteq \Pi_2^p$.

Proof. Let $L \in \text{AM}$. We note that there exists a probabilistic polynomial-time verifier V such that if $x \in L$, then for any string r that Arthur selects, there exists a string m that Merlin selects satisfying:

$$\Pr[V(x, r, m) = 1] = 1.$$

In particular, the following relation holds:

$$\forall r \exists m V(x, r, m) = 1.$$

So $L \in \Pi_2^p$. □

Remark 178. As $\text{MA}[2] \subseteq \text{AM}$, it follows immediately that $\text{MA}[2] \subseteq \Pi_2^p$. We next establish that $\text{MA}[2] \subseteq \Sigma_2^p$, which places $\text{MA}[2] \subseteq \Sigma_2^p \cap \Pi_2^p$. As AM is not known to be contained in Σ_2^p , this provides evidence that $\text{MA}[2] \subsetneq \text{AM}$.

Theorem 179. $\text{MA}[2] \subseteq \Sigma_2^p$.

Proof. Exercise. □

We conclude by showing that **Graph Isomorphism** is unlikely to be NP-complete. The key part of the proof lies in showing that if $\text{coNP} \subseteq \text{AM}$, then $\text{PH} = \Sigma_2^p \cap \Pi_2^p = \text{AM}$. It follows that if **Graph Isomorphism** is NP-complete, then **Graph Non-Isomorphism** is coNP-complete. As **Graph Non-Isomorphism** $\in \text{AM}$ (see Exercise 66), it follows that $\text{coNP} \subseteq \text{AM}$. Hence, PH collapses.

Theorem 180. Suppose $\text{coNP} \subseteq \text{AM}$. Then $\text{PH} = \Sigma_2^p \cap \Pi_2^p = \text{AM}$.

Proof. Let $L \in \Sigma_2^p$. We show that $L \in \text{AM}$. As $\text{AM} \subseteq \Pi_2^p$, this places $L \in \Pi_2^p$. As $L \in \Sigma_2^p$, there exists a verifier V such that for every $x \in L$, the following condition holds:

$$\exists y \forall z V(x, y, z) = 1. \tag{35}$$

So there exists a language $L_1 \in \text{coNP}$ such that:

$$L = \{x : \exists y, \langle x, y \rangle \in L_1\}.$$

We note that the strings z (as in (35)) serve as the coNP certificates for L_1 . As $\text{coNP} \subseteq \text{AM}$, there exists a two-round AM protocol for L_1 . We give an $\text{MA}[3]$ protocol for L as follows. Merlin begins by sending y . We then run the two-round AM protocol for L_1 . As $\text{MA}[3] = \text{AM}$, we have that $L \in \text{AM} \subseteq \Pi_2^p$. So $\Sigma_2^p \subseteq \Pi_2^p$.

It remains to show that if $\Sigma_2^p \subseteq \Pi_2^p$, then $\Sigma_2^p = \Pi_2^p$. We leave this as an exercise (see Exercise 68). It follows that $\text{PH} = \Sigma_2^p \cap \Pi_2^p = \text{AM}$. □

Corollary 181. If **Graph Isomorphism** is NP-complete, then $\text{PH} = \Sigma_2^p \cap \Pi_2^p = \text{AM}$.

Proof. Suppose that **Graph Isomorphism** is NP-complete. Then **Graph Non-Isomorphism** is coNP-complete. As **Graph Non-Isomorphism** $\in \text{AM}$ (see Exercise 66), we have that $\text{coNP} \subseteq \text{AM}$. So by Theorem 68, we have that $\text{PH} = \Sigma_2^p \cap \Pi_2^p = \text{AM}$, as desired. □

4.5.1 Exercises

(Recommended) Problem 67. Prove Theorem 179: show that $\text{MA}[2] \subseteq \Sigma_2^p$. The facts that $\text{MA}[2] \subseteq \Sigma_2^p \cap \Pi_2^p$ and AM is not known to be contained in Σ_2^p serve as evidence that $\text{MA}[2] \neq \text{AM}$.

(Recommended) Problem 68. In Theorem 180, we used the fact that if $\Sigma_2^p \subseteq \Pi_2^p$, then $\text{PH} = \Sigma_2^p \cap \Pi_2^p$. Denote $\Delta_i^p := \Sigma_i^p \cap \Pi_i^p$. Show that if $\Delta_2^p = \Sigma_2^p$, then $\Sigma_2^p = \Pi_2^p$. [**Hint:** This is analogous to the proof that if $\text{P} = \text{NP}$, then $\text{NP} = \text{coNP}$.]

5 Circuit Complexity: Razborov-Smolensky

We return to the setting of Boolean circuits to discuss Circuit Complexity. Here, we have two goals. The first is to better understand the structure between AC^0 and TC^0 . To this end, we analyze the complexity of the Parity function. The second goal is to examine the power of advice.

5.1 Razborov-Smolensky: Introduction

We refer back to Definitions 32 and 33 for the complexity classes NC^k and AC^k . While $NC^0 \subseteq AC^0$, we showed in Problem 19 that this containment is strict; namely, $NC^0 \subsetneq AC^0$. Our goal now is to understand the space between AC^0 and NC^1 . To this end, we introduce some new complexity classes.

Definition 182. Fix $k \in \mathbb{N}$ and $m \geq 2$ to be an integer. We say that $L \in AC^k[m]$ if there exists a (uniform) family of AC^k circuits $(C_n)_{n \in \mathbb{N}}$ over the basis $\{\text{AND}, \text{OR}, \text{NOT}, \text{MOD-}m\}$, where the MOD- m gates have unbounded fan-in.

Remark 183. Clearly, we have that $AC^k \subseteq AC^k[m]$.

Definition 184. Let $k \in \mathbb{N}$. We say that $L \in ACC^k$ if there exists a (uniform) family of AC^k circuits $(C_n)_{n \in \mathbb{N}}$ over the basis:

$$\{\text{AND}, \text{OR}, \text{NOT}\} \cup \{\text{MOD-}m : m \geq 2\}.$$

where the MOD- m gates have unbounded fan-in.

We introduce one final complexity class. Recall that the Threshold functions (see Exercise 16) are functions of the form:

$$\tau_t^{(n)}(x_1, \dots, x_n) = \begin{cases} 1 & \text{if } \sum_{i=1}^n x_i \geq t, \\ 0 & \text{otherwise,} \end{cases}$$

where $n \in \mathbb{N}$ and $1 \leq t \leq n$.

We now define the complexity class TC^k .

Definition 185. Let $k \in \mathbb{N}$. We say that a language $L \in TC^k$ if there exists a uniform family of circuits $(C_n)_{n \in \mathbb{N}}$ over the standard basis $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}, \text{Threshold}\}$ such that the following conditions hold:

- (a) A string $\omega \in \{0, 1\}^*$ of length n is in L if and only if $C_n(\omega) = 1$ (that is, C_n on input ω evaluates to 1).
- (b) Each circuit has fan-in 1 for the NOT gates.
- (c) The AND, OR, and Threshold gates have unbounded fan-in.
- (d) The circuit C_n has depth $O(\log^k(n))$ and size $O(n^m)$. The implicit constants and the exponent m both depend on L .

Remark 186. We now ask as to the relation between ACC^k and TC^k . It is possible to compute MOD- m gates using Threshold functions (see for instance, Problem 18). This yields the following inclusion.

Theorem 187. For each $k \in \mathbb{N}$, $ACC^k \subseteq TC^k$.

Remark 188. It remains open as to whether $ACC^k \subsetneq TC^k$ for any $k \geq 0$. In particular, we stress that it is not known whether $ACC^0 \subsetneq TC^0$.

It is also known that $TC^k \subseteq NC^{k+1}$, with $TC^0 \subsetneq NC^1$ being the only inclusion known to be strict. We record this theorem, albeit without proof. For a full proof of the fact that $TC^k \subseteq NC^{k+1}$, we defer to [Sav97, Theorem 2.11.1].

Theorem 189. For each $k \in \mathbb{N}$, we have that $TC^k \subseteq NC^{k+1}$. In particular, we have that $TC^0 \subsetneq NC^1$.

Our main goal in this section is to show that Parity $\notin AC^0$. More generally, it is known that Parity $\notin AC^0[p]$ for any prime $p \geq 3$. As Parity $\in AC^0[2]$, we have that Parity $\in ACC^0$. So we have that for any prime $p \geq 3$:

$$AC^0 \subseteq AC^0[p] \subsetneq ACC^0.$$

This result follows from the more general Razborov-Smolensky Theorem, which states that $AC^0[p]$ and $AC^0[q]$ are incomparable.

Theorem 190 (Razborov-Smolensky). Let p, q be distinct primes. Then $AC^0[p] \not\subseteq AC^0[q]$.

The proof of Theorem 207 is broken into two parts. We first show that every $AC^0[p]$ circuit on n variables can be approximated by low-degree polynomials in $\mathbb{F}_p[x_1, \dots, x_n]$. Next, we show that the MOD- q gate cannot be approximated by such low-degree polynomials.

We first arithmetize the AND, OR, and MOD- p gates as functions over \mathbb{F}_p . Namely, we have that:

$$\begin{aligned} \text{AND}(x_1, \dots, x_n) &= \prod_{i=1}^n x_i, \\ \text{OR}(x_1, \dots, x_n) &= 1 - \prod_{i=1}^n (1 - x_i), \\ \text{MOD-}p(x_1, \dots, x_n) &= 1 - (x_1 + x_2 + \dots + x_n)^{p-1}. \end{aligned}$$

Remark 191. Note that the multiplicative group \mathbb{F}_p^* has $p - 1$ elements. So if $x_1 + x_2 + \dots + x_n \equiv 0 \pmod{p}$, then $(x_1 + x_2 + \dots + x_n)^{p-1} \equiv 0 \pmod{p}$. Otherwise, one may apply Lagrange's Theorem from Group Theory (which states that in a finite group G , any element $g \in G$ satisfies that $g^{|G|} = 1$), or if one prefers- Fermat's Little Theorem, we have that $(x_1 + x_2 + \dots + x_n)^{p-1} \equiv 1 \pmod{p}$.

5.2 Razborov-Smolensky: Approximating $\text{AC}^0[p]$ Circuits

Observe that while MOD- p is always a degree $p - 1$ polynomial, regardless of the number of inputs, both AND and OR are degree n polynomials, where n is the number of inputs. Our first goal is to show how to approximate AND and OR using low-degree polynomials. We begin with the following lemma.

Lemma 192. Let $x \in \{0, 1\}^n$ is non-zero. We have that:

$$\Pr_{S \sim \text{Unif}(2^{[n]})} \left[\sum_{i \in S} x_i \neq 0 \right] \geq 1/2.$$

Here, the sum is taken over \mathbb{F}_p .

Proof. Exercise. □

Using Lemma 192, we can approximate the OR and AND function. Our approach is to use the Probabilistic Method to show the existence of such a polynomial. The idea behind the Probabilistic Method is to impose a probability measure on our ambient space. The goal is to show our desired set of objects S occurs with non-zero probability, which implies that such objects exist. We note that $\Pr[S] = 0$ does not imply the non-existence of our objects, as there can be non-empty sets of measure 0. Constructing measure 0 sets is technical; we defer the reader to standard texts on Measure Theory.

We also note that the Probabilistic Method is non-constructive. While it allows us to determine that our objects exist, it can be challenging to construct such objects explicitly.

Lemma 193. For any $m \geq 1$, there exists a polynomial $f(x) \in \mathbb{F}_p[x_1, \dots, x_n]$ of degree $(p - 1)d$ such that:

$$\Pr_{\alpha \in \{0, 1\}^n} [\text{OR}(x) \neq p(x)] \leq 2^{-m}.$$

Proof. Let $S_1, \dots, S_m \sim \text{Unif}(2^{[n]})$ by independent and identically distributed (i.i.d.). Consider the corresponding random polynomial:

$$f(x_1, \dots, x_n) := 1 - \prod_{i=1}^m \left(1 - \sum_{j \in S_i} x_j \right)^{p-1},$$

which has degree $(p - 1)m$. We note that if each $x_i = 0$, then $f(x_1, \dots, x_n) = \text{OR}(x_1, \dots, x_n)$ with probability 1. Now suppose that $x \in \{0, 1\}^n$ is non-zero. By Lemma 192, we have that for any fixed $i \in [m]$:

$$\left(1 - \sum_{j \in S_i} x_j \right)^{p-1} \equiv 0 \pmod{p}$$

with probability at most $1/2$. As S_1, \dots, S_d are i.i.d., we have that:

$$\prod_{i=1}^m \left(1 - \sum_{j \in S_i} x_j \right)^{p-1} \equiv 0 \pmod{p}$$

with probability 2^{-m} . So $f(x_1, \dots, x_n) = \text{OR}(x_1, \dots, x_n)$ with probability at least $1 - 2^{-m}$, as desired. □

We may state the analogous result for AND(x_1, \dots, x_n), which we leave as an exercise.

Lemma 194. For any $m \geq 1$, there exists a polynomial $f(x) \in \mathbb{F}_p[x_1, \dots, x_n]$ of degree $(p - 1)m$ such that:

$$\Pr_{x \in \{0, 1\}^n} [\text{AND}(x) \neq f(x)] \leq 2^{-m}.$$

We now have the technical tools to approximate $\text{AC}^0[p]$ circuits with low degree \mathbb{F}_p -polynomials. Namely, we arithmetize each gate, using approximations for the AND and OR functions. If a gate G has inputs G_1, \dots, G_k , then the polynomial G' corresponding to G is evaluated on $G'(G'_1, \dots, G'_k)$, where each G'_i is the polynomial corresponding to G_i . As our circuits have constant depth, we obtain a polynomial approximation for each circuit. We leave the precise details as an exercise.

Theorem 195. Let p be prime. Let $C(x_1, \dots, x_n)$ be an $\text{AC}^0[p]$ of depth m and size S . Then for any $m \geq 1$, there exists a polynomial $f(x) \in \mathbb{F}_p[x_1, \dots, x_n]$ of degree $((p - 1)m)^d$ such that:

$$\Pr_{x \in \{0, 1\}^n} [f(x) \neq C(x)] \leq S2^{-m}.$$

5.2.1 Exercises

(Recommended) Problem 69. Prove Lemma 192. [**Hint:** Consider the subsets $\text{supp}(x)$ that have size divisible by p . Here, $\text{supp}(x)$ is the set of indices i for which $x_i \neq 0$.]

(Recommended) Problem 70. Prove Lemma 194. [**Hint:** Write $\text{AND}(x_1, \dots, x_n)$ in terms of $\text{OR}(x_1, \dots, x_n)$. It may be helpful to start by doing so using DeMorgan's Law over the logical operators, and then to arithmetize at the end.]

(Recommended) Problem 71. Prove Theorem 195.

5.3 Razborov-Smolensky: Hilbert Functions

In the previous section, we established that every $\text{AC}^0[p]$ circuit is approximated by a low-degree polynomial over \mathbb{F}_p . Our goal now is to show that **Parity** is not approximated by such low degree polynomials. To this end, we utilize an invariant known as the (affine) Hilbert function. We note that Hilbert functions appear in Commutative Algebra as a means of measuring the complexities of graded rings. To this end, we recall some notions from Commutative Algebra. Our exposition here is adapted from [Fil10].

Definition 196. Let R be a commutative ring with identity, and let $\mathcal{I} \subseteq \mathbb{Z}$ be our indexing set. We say that R is *graded* if R admits a direct sum decomposition:

$$R = \bigoplus_{i \in \mathcal{I}} R_i$$

as Abelian groups; such that for all $i, j \in \mathcal{I}$, $R_i R_j \subseteq R_{i+j}$.

Example 197. Let $R = \mathbb{F}[x_1, \dots, x_n]$. We have that R is a graded ring over $\mathcal{I} := \mathbb{N}$, where $R_i \subseteq R$ is the set of degree i polynomials. Note that R_i is an Abelian group, but not a ring itself. However, multiplying a degree i polynomial and a degree j polynomial gives us a degree $i + j$ polynomial. That is, $R_i R_j \subseteq R_{i+j}$.

Remark 198. For our setting, we are concerned with polynomial rings over fields. We note that polynomial rings over fields are simultaneously rings and vector spaces. With this in mind, we turn to defining the Hilbert function in this restricted setting.

Definition 199. Let $R = \mathbb{F}[x_1, \dots, x_n]$ be a polynomial ring, with the grading:

$$R = \bigoplus_{i \in \mathcal{I}} R_i.$$

The *Hilbert Function* $HF_R : \mathbb{N} \rightarrow \mathbb{N}$ maps:

$$HF_R(n) = \dim_{\mathbb{F}}(R_n).$$

Remark 200. In order to reason about Boolean circuits, we will restrict attention to Hilbert functions on subsets of the Boolean cube $\{0, 1\}^n$. Let $S \subseteq \{0, 1\}^n$, and let \mathbb{F} be an arbitrary (not necessarily finite) field. Let:

$$R := \mathbb{F}[x_1, \dots, x_n] / (x_1^2 = x_1, \dots, x_n^2 = x_n),$$

which restricts our attention to the Boolean cube. For each $P \in R$, the image $P|_S$ can be viewed as a vector of length $|S|$. We consider R as a graded ring, where the grading is by degree (that is, R_i is the group of degree i polynomials in R). We are interested in examining polynomials of degree at most k over a given set $S \subseteq \{0, 1\}^n$. To this end, we denote the restricted partial sum of HF_R to be:

$$h_k(S) := \dim\{P|_S : P \in R \text{ and } \deg(P) \leq k\}.$$

As our vectors have length $|S|$, we have immediately that $h_k(S) \leq |S|$. Now we have divided out by the ideal:

$$(x_1^2 = x_1, \dots, x_n^2 = x_n),$$

So we note that each equivalence class in R has a representative, where each monomial uses only distinct variables. Thus, we also have that:

$$h_k(S) \leq \sum_{m=0}^k \binom{n}{m}.$$

Remark 201. We will see shortly that if k is small relative to $|S|$, then $h_k(S) \leq |S|/2$. This motivates the notion of the Hilbert excess function.

Definition 202. Let $S \subseteq \{0, 1\}^n$. The *Hilbert excess function* $\alpha_k(S)$ is:

$$\alpha_k(S) := 2h_k(S) - |S|.$$

Our approach to establish lower bounds is as follows:

- (a) Let $p(x_1, \dots, x_n) \in \mathbb{F}[x_1, \dots, x_n]/(x_1^2 = x_1, \dots, x_n^2 = x_n)$ be a polynomial. We may associate to $p(x)$ a Boolean function $\varphi_p(x_1, \dots, x_n)$ as follows:

$$\varphi(x_1, \dots, x_n) = \begin{cases} 1 & : p(x_1, \dots, x_n) = 0, \\ 0 & : \text{otherwise.} \end{cases}$$

Denote $Z(p(x_1, \dots, x_n))$ to be the set of points upon which $p(x_1, \dots, x_n)$ vanishes. Observe that:

$$Z(p(x)) = \overline{Z(\varphi_p(x_1, \dots, x_n))}.$$

- (b) Suppose that $\varphi(x_1, \dots, x_n)$ and $\tilde{p}(x_1, \dots, x_n)$ is a polynomial that (we hope) approximates φ . We first show that if $\tilde{p}(x_1, \dots, x_n)$ has low degree, then the complement $\overline{Z(\tilde{p}(x))} = Z(\varphi_{\tilde{p}}(x_1, \dots, x_n))$ has small excess.
- (c) Next, we show that the zero-set of a hard function (e.g., Parity, Mod- q) has high excess. To do so, we explicitly compute its Hilbert function.
- (d) As the corresponding zero sets are far apart, the low degree polynomial $\tilde{p}(x)$ is not a good approximation of our hard function.

In order to show the zero sets are far apart, it is important to understand how the Hilbert function $h_k(S)$ changes when points are added or removed. To this end, we introduce the following lemma.

Lemma 203. Let $S, T \subseteq \{0, 1\}^n$. We have that: $|\alpha_k(S) - \alpha_k(T)| \leq |S \Delta T|$.

Proof. Exercise. □

We now turn to showing that the non-zero set of a low-degree polynomial has small excess.

Theorem 204. Let $p(x)$ be a polynomial of degree d over the ring

$$R := \mathbb{F}[x_1, \dots, x_n]/(x_1^2 = x_1, \dots, x_n^2 = x_n).$$

Let $k < (n - d)/2$. Let $S = \{x : p(x) \neq 0\}$. Then $\alpha_k(S) \leq 0$.

Proof. Suppose to the contrary that $\alpha_k(S) > 0$; that is, suppose that $2h_k(S) > |S|$. We recall that $h_k(S)$ is:

$$\dim\{f|_S : f \in R \text{ and } \deg(f) \leq k\}.$$

We order the monomials of R that have degree at most k in lexicographic order. Let f_i be the i th monomial. We define a matrix M where the i th row of M is the vector $f_i|_S$. Note that $\text{rank}(M) = h_k(S)$. As $h_k(S) > |S|/2$ (using the fact that $2h_k(S) > |S|$), there exists a set $U \subseteq S$ of $h_k(S)$ linearly independent columns. So:

$$h_k(S) = h_k(U) = |U|.$$

As $h_k(S) > 2|S|$, we have that $h_k(S) > |S| - |U| = |S \setminus U|$. So there exists a non-zero polynomial $f(x)$ of degree at most k such that $f(x) = 0$ on $S \setminus U$. Our goal is to extend $f(x)$ to an indicator function on $\{0, 1\}^n$. We do this in two stages: first, we extend $f(x)$ to an indicator function on S . Then we extend $f(x)$ to an indicator function on the entire Boolean cube $\{0, 1\}^n$. We begin with some notation. For a set $T \subseteq \{0, 1\}^n$ and a point $x \in T$, denote $\delta(T, x)$ to be the elements of R that are non-zero on x and zero on the rest of T .

- **Step 1:** We begin by extending $f(x)$ to an indicator function on S . Note that as $f(x) = 0$ on $S \setminus U$, it suffices to multiply $f(x)$ by the appropriate indicator function on U . As $h_k(U) = |U|$, any Boolean function on U can be realized by a polynomial of degree at most k . As $f(x)$ is not the zero polynomial, there exists $x_0 \in U$ such that $f(x_0) \neq 0$. Now there is a non-zero polynomial $g(x) \in \delta(U, x_0)$ of degree at most k . Observe that $f(x) \cdot g(x) \in \delta(S, x_0)$.
- **Step 2:** We now extend $f(x)$ to an indicator function on the entire Boolean cube $\{0, 1\}^n$. Namely, we claim that $\tilde{f}(x) := p(x) \cdot f(x) \cdot g(x) \in \delta(\{0, 1\}^n, x_0)$. We leave the remainder of this step as an exercise.

Now observe that:

$$\begin{aligned} \deg(\tilde{f}) &= \deg(p) + \deg(f) + \deg(g) \\ &\leq d + k + k \\ &< d + 2(n - d)/2 \\ &< n. \end{aligned}$$

However, if a polynomial is zero on all but one point of $\{0, 1\}^n$, then that polynomial must have degree at least n , a contradiction. The result follows. \square

We obtain the following corollary.

Corollary 205. Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$ be a Boolean function, with the associated polynomial $p_f(x)$. Let $p(x)$ be a polynomial of degree d over the ring

$$R := \mathbb{F}[x_1, \dots, x_n]/(x_1^2 = x_1, \dots, x_n^2 = x_n).$$

Now let φ_p be the Boolean function associated with $p(x)$. Let $S := \overline{p^{-1}(0)}$. Then:

$$\Pr[f(x) \neq \varphi_p(x)] \geq 2^{-n} \cdot \alpha_{(n-d-1)/2}(S).$$

Proof. Let $k := (n - d - 1)/2$. By Theorem 204, we have that $\alpha_k(\overline{p^{-1}(0)}) \leq 0$. It follows that:

$$\alpha_k(S) - \alpha_k(\overline{p^{-1}(0)}) \geq \alpha_k(S).$$

By Lemma 203, we have that:

$$\alpha_k(S) - \alpha_k(\overline{p^{-1}(0)}) \leq |S \Delta \overline{p^{-1}(0)}|.$$

Putting the pieces together, we have that:

$$|S \Delta \overline{p^{-1}(0)}| \geq \alpha_k(S).$$

Dividing both sides by $|\{0, 1\}^n| = 2^n$, we obtain the desired result:

$$\Pr[f(x) \neq \varphi_p(x)] \geq 2^{-n} \cdot \alpha_{(n-d-1)/2}(\overline{p^{-1}(0)}).$$

\square

5.3.1 Exercises

(Recommended) Problem 72. Our goal is to prove Lemma 203. Let $S, T \subseteq \{0, 1\}^n$. We have that: $|\alpha_k(S) - \alpha_k(T)| \leq |S \Delta T|$.

(a) It suffices to prove this lemma in the case when $T = S \cup \{x\}$ for some $x \notin S$. Start by showing that:

$$h_k(S) \leq h_k(T) \leq h_k(S) + 1.$$

(b) Deduce that $\alpha_k(S) - 1 \leq \alpha_k(T) \leq \alpha_k(S) + 1$.

(Recommended) Problem 73. Complete Step 2 in the proof of Theorem 204 by showing that $\tilde{f}(x) \in \delta(\{0, 1\}^n, x_0)$.

(Recommended) Problem 74. In light of Corollary 205, prove the following. Suppose that $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a Boolean function, with the associated polynomial $p_f(x)$. Let $S := \overline{p_f^{-1}(0)}$. Suppose that for some $k < n/2$, we have that:

$$h_k(S) = \sum_{i=0}^k \binom{n}{i}.$$

Let $p(x)$ be a polynomial of degree less than $n - 2k$, and let φ_p be the associated Boolean function. We have that:

$$\Pr[f(x) \neq \varphi_p(x)] \geq 2\Pr[\text{Binom}(n, 1/2) \leq k] - \Pr[f = 0].$$

5.4 Razborov-Smolensky: Obtaining Circuit Lower Bounds

We begin with the following lemma.

Lemma 206. Let $S \subseteq \{0, 1\}^n$. Then for any k , we have that:

$$h_k(S) = \sum_{i=0}^k \binom{n}{i}$$

if and only if the only polynomial of degree at most k which is zero on all of S is the zero polynomial.

Proof. We note that $h_k(S)$ is the dimension of the subspace $V := \{P|_S\}$ of polynomials of degree at most k over the ring

$$R = \mathbb{F}[x_1, \dots, x_n]/(x_1^2 = x_1, \dots, x_n^2 = x_n).$$

We note that the monomials of degree at most k span V . The number of such monomials is:

$$\sum_{i=0}^k \binom{n}{i}.$$

So the monomials of degree at most k form a basis for V if and only if they are linearly independent over S . The result follows. \square

Theorem 207. Let \mathbb{F} be a field, and let q be a prime different from $\text{char}(\mathbb{F})$. Let S be the set of vectors whose Hamming weights are not multiples of q . For all $k < n/2$, we have that:

$$h_k(S) = \sum_{i=0}^k \binom{n}{i}.$$

Proof. We note that an \mathbb{F} -vector space V can be viewed as a K -vector space, where K/\mathbb{F} is a field extension, then $\dim_K(V) \leq \dim_{\mathbb{F}}(V)$. So without loss of generality, suppose that \mathbb{F} is algebraically closed (otherwise, we consider $K = \overline{\mathbb{F}}$, where $\overline{\mathbb{F}}$ is the algebraic closure of \mathbb{F}). Now observe that $x^q - 1$ and its derivative qx^{q-1} are relatively prime. Thus, $x^q - 1$ does not have any repeated roots. As \mathbb{F} is algebraically closed, all solutions to $x^q - 1$ are contained in \mathbb{F} . In particular, there is a primitive q th root of unity $\omega \in \mathbb{F}$.

By Lemma 206, it suffices to show that the only polynomial that is zero on all of S is the zero polynomial. Let P be a polynomial that is zero on S . We apply the substitution $\hat{f} : x_i \mapsto (\omega - 1)x_i + 1$, where we denote $y_i := (\omega - 1)x_i + 1$, to obtain a new polynomial $P_y(y_1, \dots, y_n)$ of the same degree. Observe that $x_i = 0$ corresponds to $y_i = 1$, and that $x_i = 1$ corresponds to $y_i = \omega$. We also note that this transformation admits an inverse. Namely:

$$y_i \mapsto \frac{y_i - 1}{\omega - 1}.$$

So P is zero on S if and only if P_y is zero on:

$$\begin{aligned} S_y &= \{y \in \{1, \omega\}^n : y = \hat{f}(x) \text{ and } x \in S\} \\ &= \{y \in \{1, \omega\}^n : y_1 \cdots y_n \neq 1\}. \end{aligned}$$

Now define the polynomial:

$$Q = P_y(1 - y_1 \cdots y_n).$$

If $(y_1, \dots, y_n) \in S_y$, then $P_y = 0$. Otherwise, $y_1 \cdots y_n = 1$; in which case,

$$Q = P_y - P_y = 0.$$

So Q is zero on $\{1, \omega\}^n$. Let $U \subseteq [n]$, and let $M_U = \prod_{i \in U} y_i$ be the corresponding monomial. For each $U \subseteq [n]$, we have that:

$$y_1 \cdots y_m M_U = c \prod_{h \notin U} y_h + \cdots,$$

where \cdots on the RHS represents monomials of higher degree. We apply use the identity that $y_i^2 = (\omega + 1)y_i - \omega$ (see Exercise 76) to obtain a new, square-free polynomial R that agrees with Q and is zero on $\{1, \omega\}^n$. As the

square-free monomials are linearly independent, we have that $R = 0$.

We now claim that $P_y = 0$. Suppose to the contrary that $P_y \neq 0$. Let M be a monomial of maximal degree in P_y , corresponding to the set $U \subseteq [n]$. As P_y has degree $k < n/2$, we have that $M_{\overline{U}} \in R$, which contradicts the fact that $R = 0$. Thus, $P_y = 0$. So $P = 0$. \square

We now put the pieces together to obtain our circuit lower bounds. Namely, we show that asymptotically, low-degree polynomials do not approximate the MOD- q function any better than random noise.

Theorem 208. Let p be prime, and let C be an $\text{AC}^0[p]$ of size S and depth h . Let $f = \text{MOD-}q$, for some prime $q \neq p$. Then:

$$\Pr[C(x) = f(x)] = 1/2 + O\left(\frac{(\log(nS))^h}{\sqrt{n}}\right) + \frac{1}{n}.$$

Proof. Let $m = \log(nS)$. By Theorem 195, there exists an \mathbb{F}_p polynomial $p(x)$ of degree $d := ((p-1)m)^h$ that differs from C with probability at most $S2^{-m} = 1/n$. By Theorem 207 and Exercise 74, we have that:

$$\Pr[C \neq f] \geq 2\Pr[\text{Binom}(n, 1/2) \leq n/2 - t] - 1/2 - 1/n,$$

where:

$$t := \frac{n}{2} - \frac{n-d+1}{2} = O((\log(nS))^h).$$

Now we estimate each binomial coefficient with the central one to obtain:

$$\Pr[\text{Binom}(n, 1/2) \leq n/2 - t] = \frac{1}{2} - 2^{-n} \sum_{i=0}^{t-1} \binom{n}{i} = \frac{1}{2} - O\left(\frac{t}{\sqrt{n}}\right).$$

Thus:

$$\begin{aligned} \Pr[C \neq f] &\geq 2\Pr[\text{Binom}(n, 1/2) \leq n/2 - t] - 1/2 - 1/n \\ &= 2\left[\frac{1}{2} - O\left(\frac{t}{\sqrt{n}}\right)\right] - 1/2 - 1/n \\ &= 1/2 - O\left(\frac{t}{\sqrt{n}}\right) - 1/n. \end{aligned}$$

\square

5.4.1 Exercises

(Recommended) Problem 75. Let \mathbb{F} be a field, and let q be a prime different from $\text{char}(\mathbb{F})$. Let S be the set of vectors whose Hamming weights are not multiples of q . Let ω be a primitive q th root of unity. Suppose that $P(x_1, \dots, x_n) \in \mathbb{F}[x_1, \dots, x_n]$ is zero on S . Consider the substitution $\hat{f} : x_i \mapsto (\omega - 1)x_i + 1$, where we denote $y_i := (\omega - 1)x_i + 1$, to obtain a new polynomial $P_y(y_1, \dots, y_n)$ of the same degree. Show that P_y is zero on the set:

$$S_y = \{y \in \{1, \omega\}^n : y_1 \cdots y_n \neq 1\}.$$

Hint: It may be helpful to take $q = 3$ and apply the transformation to $x_1x_2x_3$.

(Recommended) Problem 76. Let \mathbb{F} be a field, and let q be a prime different from $\text{char}(\mathbb{F})$. Let S be the set of vectors whose Hamming weights are not multiples of q . Let ω be a primitive q th root of unity. Let $x \in \{0, 1\}$. Observe that $x^2 = x$. Now define the transformation $x \mapsto (\omega - 1)x + 1$, where $y = (\omega - 1)x + 1$. Show that:

$$y^2 = (\omega + 1)y - \omega.$$

Hint: Evaluate both sides of the equation separately.

(Recommended) Problem 77. Let p, q be distinct primes. We have shown that $\text{MOD-}q \notin \text{AC}^0[p]$. Deduce that $\text{AC}^0 \subsetneq \text{AC}^0[p]$.

(Recommended) Problem 78. Show that $\text{AC}^0[2]$ and $\text{AC}^0[3]$ are both strictly contained in $\text{AC}^0[6]$.

6 Circuit Complexity: The Power of Advice

In this section, we explore the power of advice. Here, advice is closely related to the notion of a certificate. We motivate it as follows. Suppose we have a trustworthy Instructor that is all-knowing and wants to help students solve problems. While the Instructor would be able to solve the problem if they looked at it, they are often busy and instead give the same advice to all students. In particular, for students trying to decide whether the string ω belongs to the language L , the Instructor gives advice that depends only on L and $|\omega|$. We first aim to introduce complexity classes related to advice, as well as to understand their structural properties.

6.1 Computation with Advice

We begin by introducing the notion of advice for Turing Machines.

Definition 209. Let $T, a : \mathbb{N} \rightarrow \mathbb{N}$ be functions. We say that the language $L \in \text{DTIME}(T(n))/a(n)$ precisely if there exists a deterministic Turing Machine M and a sequence of strings $(\alpha_n)_{n \in \mathbb{N}}$ with $\alpha_n \in \{0, 1\}^{a(n)}$ such that:

$$M(x, \alpha_{|x|}) = 1 \iff x \in L$$

and $M(x, \alpha_{|x|})$ runs in time $O(T(|x|))$.

The analogous definition holds for non-deterministic computation.

Definition 210. Let $T, a : \mathbb{N} \rightarrow \mathbb{N}$ be functions. We say that the language $L \in \text{NTIME}(T(n))/a(n)$ precisely if there exists a non-deterministic Turing Machine M and a sequence of strings $(\alpha_n)_{n \in \mathbb{N}}$ with $\alpha_n \in \{0, 1\}^{a(n)}$ such that:

$$M(x, \alpha_{|x|}) = 1 \iff x \in L$$

and $M(x, \alpha_{|x|})$ runs in time $O(T(|x|))$.

This leads us to our definition of P/poly , which are languages decidable by polynomial-time deterministic Turing Machines that take polynomial advice strings.

Definition 211 (P/poly). Define:

$$\text{P/poly} := \bigcup_{c, d \in \mathbb{N}} \text{DTIME}(n^c)/n^d.$$

Remark 212. We note that polynomial-time deterministic Turing Machines that don't take advice capture precisely P . That is, $\text{P} = \text{P}/0$. Allowing for even a constant amount of advice already increases the power beyond P .

Theorem 213. $\text{P} \subsetneq \text{P}/1$.

Proof. Clearly $\text{P} \subseteq \text{P}/1$. We note that P contains only decidable problems. So in order to separate P from $\text{P}/1$, we show that $\text{P}/1$ contains an undecidable problem. Recall:

$$L_{\text{diag}} = \{\omega_i : \omega_i \text{ is the } i\text{th string in } \Sigma^*, \text{ which is accepted by the } i\text{th Turing Machine } M_i\},$$

which is undecidable. Now define:

$$L' := \{1^i : \omega_i \in L_{\text{diag}}\}.$$

In particular, $L_{\text{diag}} \leq_m L'$, so L' is undecidable. We now claim that $L' \in \text{P}/1$. Consider the sequence of advice strings $(\alpha_i)_{i \in \mathbb{N}}$, where:

$$\alpha_i = \begin{cases} 1 & : 1^i \in L', \\ 0 & : \text{otherwise.} \end{cases}$$

Let M be a $\text{P}/1$ -TM, and let $\omega \in \{1\}^*$, where M is defined such that $M(\omega, \alpha_{|\omega|}) = 1$ if and only if $\alpha_{|\omega|} = 1$. However, $\alpha_{|\omega|} = 1$ if and only if $\omega \in L'$. Thus, M decides L' . It follows that $L' \in \text{P}/1$. As $L' \notin \text{P}$, we have that $\text{P} \subsetneq \text{P}/1$, as desired. \square

It is natural to ask as to the power of P/poly; in particular, whether $P/1 = P/\text{poly}$. This is a question perhaps best answered in the realm of circuits. We note that Boolean circuits are logspace equivalent to Turing Machines. Savage covers the proof in depth (see [Sav97, Chapter 3]). We briefly sketch the approach for simulating circuits by Turing Machines. The key idea is that our Turing Machine performs a breadth-first traversal over the circuit, starting from the output node. It labels the internal gates based on their dependencies. Once these labels are in place, the Turing Machine performs a second pass of the circuit to evaluate it on the specified inputs. So effectively, the size of the circuit (which we recall is the number of gates) corresponds precisely to the serial runtime of the corresponding Turing Machine. Thus, if we want polynomial-time Turing Machines, we consider circuits with a polynomial-number of gates. This motivates the following complexity class.

Definition 214. Let $T : \mathbb{N} \rightarrow \mathbb{N}$. We say that $L \in \text{SIZE}(T(n))$ if there exists a family of Boolean circuits $(C_n)_{n \in \mathbb{N}}$ over $\Omega_0 = \{\text{AND}, \text{OR}, \text{NOT}\}$ satisfying the following conditions:

- (a) AND and OR have fan-in 2.
- (b) We have that $x \in L$ if and only if $C_n(x) = 1$, where $n := |x|$, and
- (c) C_n has size $T(n)$.

Our goal now is to characterize P/poly in terms of circuits. From our above discussions, we have enough to sketch a proof that:

$$\text{P/poly} = \bigcup_{k \in \mathbb{N}} \text{SIZE}(n^k).$$

Given a P/poly-TM M with the corresponding advice strings $(\alpha_n)_{n \in \mathbb{N}}$, there exists a polynomial-sized family of circuits (C_n) that simulates M . For each circuit C_n in our family, we may hardware the input bits corresponding to α_n . So:

$$\text{P/poly} \subseteq \bigcup_{k \in \mathbb{N}} \text{SIZE}(n^k).$$

To obtain the other inclusion, we take a polynomial-sized family of circuits $(C_n)_{n \in \mathbb{N}}$ and use their encodings as the advice strings $(\alpha_n)_{n \in \mathbb{N}}$. Namely, α_n is the encoding of C_n . So on input $\langle x, \alpha_n \rangle$, M uses α_n to construct a simulate C_n on x . M accepts $\langle x, \alpha_n \rangle$ if and only if $C_n(x) = 1$. So:

$$\bigcup_{k \in \mathbb{N}} \text{SIZE}(n^k) \subseteq \text{P/poly}.$$

So while we can characterize P/poly in terms of circuits, the above argument provides little insight into the structure of P/poly. Namely, it does not allow us to answer whether $P/1 = P/\text{poly}$. We delve a bit deeper. Our approach to characterize P/poly is based on [Vol99].

Theorem 215. Let $s(n) \geq n$. Then:

$$\text{SIZE}(s(n)) \subseteq \text{DTIME}((s(n))^2)/O(s \log(s)).$$

Proof. The key idea is to use a Turing Machine to simulate a family of circuits $(C_n)_{n \in \mathbb{N}}$, where C_n has size $s(n)$ and each gate has bounded fan-in. We fix an encoding scheme for \mathcal{C} , where the gates of a given circuit are topologically ordered. That is, the label of a gate g is larger than the labels for all of its predecessor gates. Let $f : \{0, 1\}^* \rightarrow \{0, 1\}$ be given by $f(x) := \langle C_{|x|} \rangle$, where $\langle C_{|x|} \rangle$ is the encoding of the circuit $C_{|x|}$ as a natural number. Note that a single gate in C_n uses at most $\lceil \log(s(n)) \rceil$ bits. So $|f(x)| = O(s(|x|) \cdot \log(s(|x|)))$.

Now that we have a way of encoding each circuit, we design a Turing Machine M to simulate the circuits. On input $\langle x, C_n \rangle$, where $n := |x|$, M simulates C_n one gate at a time in the ordering specified by the encoding. As the gates were topologically ordered, we may store the values of each gate g on the work tape as we evaluate g . Note that in order to evaluate g , we examine at most every gate considered thus far. As there are $s(n)$ gates to evaluate, simulating the circuit takes time $O((s(n))^2)$. The result follows. \square

We obtain the following corollary.

Corollary 216. Let $t(n) \geq n$. Then:

$$\text{SIZE}(t^{O(1)}) = \text{DTIME}(t^{O(1)})/t^{O(1)}.$$

Proof. By Theorem 215, we have that:

$$\text{SIZE}(t^{O(1)}) \subseteq \text{DTIME}(t^{O(1)})/t^{O(1)}.$$

Now let $L \in \text{DTIME}(t^{O(1)})/t^{O(1)}$, and let M be the $\text{DTIME}(t^{O(1)})$ -TM that accepts L with the aid of the corresponding advice function $f(n)$. We may construct an equivalent family of circuits $\mathcal{C} := (C_n)_{n \in \mathbb{N}}$ with size $t^{O(1)}$ that simulates M . In particular, we hard-wire $f(n)$ into the construction of C_n . So $L \in \text{SIZE}(t^{O(1)})$, which yields that:

$$\text{DTIME}(t^{O(1)})/t^{O(1)} \subseteq \text{SIZE}(t^{O(1)}).$$

The result follows. □

Remark 217. Corollary 216 yields that $\text{SIZE}(n^{O(1)}) = \text{P/poly}$.

We now seek to resolve whether $\text{P}/1 = \text{P/poly}$. To this end, we introduce a hierarchy theorem for SIZE . Recall Shannon's Theorem (Theorem 38), which states that for every $\epsilon \in (0, 1)$ and every $n \geq 6$, the fraction of Boolean functions $f : \{0, 1\}^n \rightarrow \{0, 1\}$ using more than:

$$\frac{2^n}{2n}(1 - \epsilon)$$

gates is at least $1 - 2^{-\epsilon \cdot 2^n}$. We obtain the Size Hierarchy Theorem as a corollary. Our proof is adapted from Mulzer [Mul].

Theorem 218 (Size Hierarchy Theorem). Suppose that $f(n) \in o(g(n))$. Then $\text{SIZE}(f(n)) \subsetneq \text{SIZE}(g(n))$.

Proof. Let k be the largest integer such that:

$$f(n) < 2^k/2k < g(n)$$

By Shannon's Theorem (Theorem 38), there exists a Boolean function $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ that can be realized by a circuit of size at most $2^k/2k$, but not by a circuit of size $f(n)$. Define:

$$L_n := \{\omega 0^{n-k} : \omega \in \{0, 1\}^k, \varphi(\omega) = 1\}.$$

Now define:

$$L = \bigcup_{n \geq 6} L_n.$$

So $L \in \text{SIZE}(g(n))$, but $L \notin \text{SIZE}(f(n))$. The result follows. □

We obtain as a corollary that $\text{P}/1 \subsetneq \text{P/poly}$.

Corollary 219. We have that $\text{P}/1 \subsetneq \text{P/poly}$.

6.1.1 Exercises

(Recommended) Problem 79. We showed in Theorem 213 that $\text{P} \subsetneq \text{P/poly}$ by showing that P/poly contains an undecidable language. In this problem, we aim to separate P from P/poly using a decidable language.

- (a) Show that there exists a language $L \in \text{DTIME}(2^{2^n})$ such that for all $k \in \mathbb{N}$, $L \notin \text{DTIME}(2^{n^k})$.
- (b) Let L be the language from part (a). Define $U_L := \{1^n : \langle n \rangle \in L\}$. Here, $\langle n \rangle$ is the binary representation of the natural number n . Show that $U_L \in \text{DTIME}(2^{n^k})$. [**Note:** We may think of any language as a subset of \mathbb{N} . So n here is a natural number.]
- (c) Show that if $U_L \in \text{P}$, then $L \in \text{DTIME}(2^{n^k})$ for some $k \in \mathbb{N}$. Conclude that $U_L \notin \text{P}$.
- (d) Show that $U_L \in \text{P/poly}$.

(Recommended) Problem 80. A language L is said to be *sparse* if there exists a polynomial $p(n)$ such that for all n :

$$|L \cap \{0, 1\}^n| \leq p(n).$$

That is, for each n , L has at most polynomially many strings of length n . Show that every sparse language is in P/poly .

(Recommended) Problem 81. Define P/\log to be the set of languages decidable in polynomial-time using advice strings of length $O(\log(n))$. Show that if $3\text{SAT} \in \text{P}/\log$, then $\text{P} = \text{NP}$. You may assume that every 3SAT formula on n variables has the same size. [**Hint:** Apply the same technique as in the search-to-decision reduction for SAT .]

6.2 Sparse Sets and Mahaney's Theorem

In this section, we examine the structure of sparse sets. Namely, we ask whether sparse sets can be NP-complete. Mahaney showed that unless $P = NP$, no sparse set is NP-hard [Mah82]. We present a proof due to Agrawal (the key technique is presented in [AA96]). Our exposition of this proof is adapted from [Gro16].

We begin by recalling the definition of a sparse set.

Definition 220. A language L is said to be *sparse* if there exists a polynomial $p(n)$ such that for all n :

$$|L \cap \{0, 1\}^n| \leq p(n).$$

Theorem 221 (Mahaney [Mah82]). If $P \neq NP$, then no sparse set is NP-hard.

Proof. The proof is by contrapositive. Suppose L is a sparse NP-hard language. We will show how to solve SAT in polynomial time. Let $\varphi(x_1, \dots, x_n)$ be a Boolean formula. Our goal is to decide whether φ is satisfiable. We consider the downward self-reduction tree, which is rooted at $\varphi(x_1, \dots, x_n)$. The left child of the root is $\varphi_0(0, x_2, \dots, x_n)$, and the right child of the root is $\varphi_1(1, x_2, \dots, x_n)$. We apply this construction recursively for $\varphi_0(0, x_2, \dots, x_n)$ and $\varphi_1(1, x_2, \dots, x_n)$.

Now observe that $\varphi(x_1, \dots, x_n)$ is satisfiable if and only if for each level ℓ of the tree, at least one formula at level ℓ is satisfiable. Using the downward reducibility tree outright takes exponential time, as it amounts to a brute-force approach over all possible 2^n assignments. Our goal instead is to prune the tree at each level, so that the following conditions hold:

- φ is satisfiable if and only if for each level ℓ , at least one formula at level ℓ is satisfiable.
- At each level ℓ , there are only polynomially many strings.
- We can prune the tree in polynomial time.

If we can modify the tree to satisfy the above conditions, then the resulting SAT algorithm would run in polynomial time (as there are only n levels).

Suppose that $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$ is a polynomial-time many-to-one reduction from SAT to L . Let $s(n)$ be the polynomial bounding the number of strings in L of length at most n . Let $r(n)$ be a polynomial bound on the length that f stretches the input strings; that is, $|f(\varphi)| \leq r(|\varphi|)$ for all φ . Note that as f is polynomial-time computable, f may only stretch the input strings by a polynomial amount. Our goal now is to prune the tree to have at most $t(n) + 1$ formulas at each level, where $t(n) = s(r(2n + 5))$, which is polynomially-bounded. Note that the $2n + 5$ comes from applying the reduction f to strings of the form $(\varphi) \vee (\psi)$. The $() \vee ()$ strings add 5 letters, while φ, ψ each have length n .

Our construction proceeds in stages, starting with the root node. Let $\varphi_1, \dots, \varphi_k$ be the formulas at the current level ℓ of the tree. If $k \leq t(n) + 1$, we continue to the next stage. Suppose instead that $k > t(n) + 1$. For each $i \in \{2, \dots, k\}$, let $q_i := f((\varphi_1) \vee (\varphi_i))$. We have the following cases:

- **Case 1:** Suppose that q_2, \dots, q_k are all distinct. We claim that φ_1 is not satisfiable. As $k > t(n) + 1$, there exists $q_i \notin L \cap \Sigma^{\leq r(2n+5)}$. Thus, as f is a reduction, $f((\varphi_1) \vee (\varphi_i)) \notin L$. So $\varphi_1 \vee \varphi_i$ is not satisfiable, which implies that φ_1 is not satisfiable. So we may safely remove φ_1 from the tree.
- **Case 2:** Suppose that $q_i = q_j$ for distinct i, j . In this case, we remove φ_i . We leave it as an exercise to justify why we can remove φ_i safely.

Note that we have only removed one formula from level ℓ . It may still be the case that $k - 1 > t(n) + 1$. We relabel the remaining formulas starting from 1 and proceed again if $k - 1 > t(n) + 1$. Otherwise, we proceed to level $\ell + 1$ and continue our pruning. The result follows. \square

We now address some fundamental questions about the P vs. NP problem with respect to sparse sets. First, it is natural to ask if $P \neq NP$, yet there is an algorithm for SAT that runs in polynomial-time on most instances. Here, we consider *most* to mean *on all but a sparse set*. The answer to this question is *no*.

Theorem 222. Suppose that $P \neq NP$. There is no algorithm that solves SAT correctly and runs in polynomial-time on all but a sparse set.

Proof. The proof is by contrapositive. Suppose we have an algorithm A that solves SAT and runs in polynomial-time on all instances, save for those in a sparse set L . Let $T(n)$ denote the runtime complexity of A on instances not in L . Without loss of generality, suppose that A takes strictly more time than $T(|x|)$ on input $x \in L$ (otherwise, we may select a smaller sparse set that does not contain x). Let $L' := L \cap \text{SAT}$. Observe that L is sparse. We show that if L' is empty, then SAT is in P. Otherwise, L' is NP-hard.

Suppose first that L' is empty. Then A fails to run in time $T(|x|)$ only for strings $x \notin \text{SAT}$. So on input x , we clock A to ensure that A does not run for more than $T(|x|)$ steps. If A fails to make a decision in $T(|x|)$ steps, then we reject x , as $x \notin \text{SAT}$. So $\text{SAT} \in P$, and we conclude that $P = NP$.

Suppose instead that L' is non-empty. So there exists a string $x_{\text{yes}} \in L'$, as well as a string $x_{\text{no}} \notin L'$. We note that x_{yes} exists, as $L' \neq \emptyset$. Now x_{no} exists, as L' is sparse. We construct a reduction $\varphi : \Sigma^* \rightarrow \Sigma^*$ from SAT to L' . On input x , we run $A(x)$ for $T(|x|)$ steps. If $A(x) = 1$, then we output x_{yes} . If instead, $A(x) = 0$, we output x_{no} . Otherwise, we output x . So $x \in \text{SAT}$ if and only if $\varphi(x) \in L'$. So if L' is non-empty, then L' is NP-hard. As L' is sparse, we have by Mahaney's Theorem (Theorem 221) that $P = NP$. This completes the proof. \square

As we cannot solve SAT correctly on all instances and in polynomial-time on all but a sparse set, it is natural to ask whether we can solve SAT correctly on all but a sparse set and in polynomial-time on all but the same sparse set. It turns out that this weaker condition also implies that $P = NP$. Note that we refer to such sets as P-close, which we formalize below.

Definition 223. A language L is said to be P-close if there exists a language $L' \in P$ such that the symmetric difference $L \Delta L'$ is sparse.

Theorem 224. Suppose that SAT is P-close. Then $P = NP$.

Proof. Suppose that SAT is P-close. So there exists a language $L \in P$ such that $S := \text{SAT} \Delta L$ is sparse. We show that S is NP-hard under Turing reductions. Let A be a polynomial-time algorithm for L . On input x , we run $A(x)$ to determine whether $x \in L$. If $A(x) = 0$, we have that $x \in \text{SAT} \iff x \in S$. If instead we have that $A(x) = 1$, then $x \in \text{SAT} \iff x \notin S$. In either case, a single oracle query to S will yield the answer.

We now modify Agrawal's proof of Mahaney's Theorem (Theorem 221) to show that if there is a sparse NP-hard set under polynomial-time Turing reductions, then $P = NP$. We think of the reduction in two parts. The first part consists of a polynomial-time computable function $f : \{0, 1\}^* \rightarrow \{0, 1\}^*$, which determines the query to make. Here, we think of $f(x)$ as first computing $A(x)$ to decide whether $x \in L$. If $A(x) = 0$, then $f(x)$ provides the query to ask if $x \in S$. Otherwise, $f(x)$ provides the query to ask if $x \notin S$. The second part consists of a polynomial-time computable function $g : \{0, 1\}^* \rightarrow \{\text{orig}, \text{neg}\}$, which tells us whether we are asking if $x \in S$ (orig) or whether $x \notin S$ (neg).

We again consider the downward self-reducibility tree rooted at φ . As before, let $r(n)$ be the polynomial-bound on the length that f stretches the input strings; that is, $|f(\varphi)| \leq r(|\varphi|)$ for all φ . Let $s(n)$ be the polynomial bound on the number of strings of length at most n that appear in S . Let $t(n) = s(r(2n+5))$. Our goal now is to prune the tree, so that there are at most $2(t(n)+1)$ formulas with at least one satisfiable formula at each level.

Our construction proceeds in stages, starting with the root node. Let $\varphi_1, \dots, \varphi_k$ be the formulas at the current level ℓ of the tree. If $k \leq 2(t(n)+1)$, we continue to the next stage. Suppose instead that $k > t(n)+1$. For each $i \in \{2, \dots, k\}$, we construct a pair of formulas $(q_i, b_i) := (f((\varphi_1) \vee (\varphi_i)), g(\varphi_1 \vee \varphi_i))$. Here, we think of q_i as the query. Now $g(\varphi_1 \vee \varphi_i)$ is telling us whether we are asking whether $f((\varphi_1) \vee (\varphi_i)) \in S$ or whether $f((\varphi_1) \vee (\varphi_i)) \notin S$.

- **Case 1:** Suppose that q_2, \dots, q_k are all distinct. Then there are at least $t(n)+1$ indices i such that $b_i = \text{orig}$, or there are at least $t(n)+1$ indices i such that $b_i = \text{neg}$. Suppose first that there are at least $t(n)+1$ indices i such that $b_i = \text{orig}$. So $\varphi_1 \vee \varphi_i \notin L$. It follows that $\varphi_1 \vee \varphi_i \in \text{SAT}$ if and only if $\varphi_1 \vee \varphi_i \in S$. As there are at least $t(n)+1$ indices i such that $b_i = \text{orig}$, there exists an index i such that

$q_i \notin S \cap \{0, 1\}^{\leq r(2n+5)}$. As $q_i = f((\varphi_1) \vee (\varphi_i)) \notin S$, we have that $\varphi_1 \vee \varphi_i \notin \text{SAT}$. So $\varphi_1 \notin \text{SAT}$. Thus, we may safely remove φ_1 .

Suppose instead that there are at least $t(n) + 1$ indices i such that $b_i = \text{neg}$. So there exists an index i such that some q_i with $b_i = \text{neg}$ does not belong to S . This implies that $\varphi_1 \vee \varphi_i$ is satisfiable. So the original formula φ is satisfiable, and we may stop.

- **Case 2:** Suppose that $q_i = q_j$ for distinct i, j . If $b_i = b_j$, then we may remove φ_i safely. Suppose instead that $b_i \neq b_j$. We claim that φ is satisfiable. WLOG, suppose that $b_i = \text{orig}$ and $b_j = \text{neg}$. Let $q := q_i = q_j$. If $q \in L$, then (q_i, b_i) yields the answer that $q_i \in \text{SAT}$. If $q \notin L$, then (q_j, b_j) yields the answer that $q_j \in \text{SAT}$. In either case, $\varphi \in \text{SAT}$.

Note that we have only removed one formula from level ℓ . It may still be the case that $k - 1 > 2(t(n) + 1)$. We relabel the remaining formulas starting from 1 and proceed again if $k - 1 > 2(t(n) + 1)$. Otherwise, we proceed to level $\ell + 1$ and continue our pruning. The result follows. \square

6.2.1 Exercises

(Recommended) Problem 82. Recall Case 2 in the proof of Mahaney's Theorem (Theorem 221). In this case, we have that $q_i = q_j$ for some distinct i, j . Carefully explain why we can safely remove φ_i from level ℓ of the tree (where level ℓ is the current level we are considering).

(Recommended) Problem 83. The Tautology problem takes as input a Boolean formula $\varphi(x_1, \dots, x_n)$ and asks whether every assignment satisfies φ . We note that Tautology is coNP-complete. Fortune's Theorem states that no coNP-hard set is sparse unless $\text{P} = \text{NP}$. Modify the proof of Mahaney's Theorem (Theorem 221) to prove Fortune's Theorem. [**Hint:** In our tree, we want *every* formula at a given level to be satisfiable, rather than just having some satisfiable formula at each level. We also want to consider $q_i := f(\varphi_1 \wedge \varphi_i)$. Why is this the case?]

6.3 Karp-Lipton Theorems

In the previous section, we showed that P/poly is quite a bit more powerful than P, in that P/poly contains undecidable languages. It is natural to ask whether $\text{NP} \subseteq \text{P/poly}$. The Karp-Lipton-Sipser Theorem asserts that this is unlikely; for if $\text{NP} \subseteq \text{P/poly}$, then $\text{PH} = \Sigma_2^p \cap \Pi_2^p$. We note that the advice strings for P/poly-TMs depend only on the length of the string ω , which we are testing to see whether ω is in the language L . If $L \in \text{NP}$, then there may be two strings $\omega_1, \omega_2 \in L$ for which we associate different certificates of the same length. For this reason, it seems unlikely that P/poly is sufficiently powerful to capture NP.

Theorem 225 (Karp-Lipton-Sipser, 1980). If $\text{NP} \subseteq \text{P/poly}$, then $\text{PH} = \Sigma_2^p \cap \Pi_2^p$.

We begin by showing that if $\text{NP} \subseteq \text{P/poly}$, the polynomial-time search-to-decision reduction for SAT can be executed by a family of polynomial-sized circuits. That is, the search-to-decision reduction is computable in P/poly.

Lemma 226. Suppose that $\text{NP} \subseteq \text{P/poly}$. Then for every polynomial-time computable function $\varphi(x, y)$, there exists a family of polynomial-sized circuits $(C_n)_{n \in \mathbb{N}}$ such that:

$$C_{|x|}(x) = \begin{cases} y & : \exists y \varphi(x, y) = 1, \\ \text{a sequence of zeros otherwise.} \end{cases}$$

Proof. Let $p(n)$ be a polynomial such that (WLOG), $|y| = p(n)$. Define the circuits $C_n^1, \dots, C_n^{p(n)}$ as follows. On input $\langle x, b_1, \dots, b_{i-1} \rangle$, where $b_1, \dots, b_{i-1} \in \{0, 1\}$, C_n^i outputs 1 if and only if there is a satisfying assignment for $\varphi(x, y) = 1$, where $y_1 = b_1, \dots, y_{i-1} = b_{i-1}$, and $y_i = 1$. Now each C_n^i realizes an NP computation and so can be realized using only a polynomial number of gates. Now $C_{|x|}(x)$ computes in sequence $C_n^1(x), C_n^2(\langle x, b_1 \rangle), \dots, C_n^{p(n)}(\langle x, b_1, \dots, b_{n-1} \rangle)$. The result follows. \square

We now prove the Karp-Lipton-Sipser Theorem.

Proof of Theorem 225. Let $L \in \Pi_2^p$. So there exists a Π_2^p -TM M such that:

$$\omega \in L \iff \forall x_1 \exists x_2 M(\omega, x_1, x_2) = 1.$$

Define:

$$L' := \{ \langle x, y_1 \rangle : \exists y_2 \text{ s.t. } M(x, y_1, y_2) = 1 \}.$$

As M runs in time $\text{poly}(|\omega|)$, we have that M is an NP machine that computes L' . Here, y_2 is the certificate associated with the pair $\langle x, y_1 \rangle$. Now as $\text{NP} \subseteq \text{P/poly}$, there exists a polynomial $p(n)$ and a family $(C_n)_{n \in \mathbb{N}}$ of circuits, where C_n has size $p(n)$, that computes L' . By Lemma 226, there exists a circuit C_n of size polynomial in n such that for every ω of length n and every x_1 of length at most polynomial in n :

$$\exists x_2 M(\omega, x_1, x_2) = 1 \iff M(\omega, x_1, C_n(\omega, x_1)) = 1.$$

Let $q(n)$ be the polynomial bound on C_n . So we have that:

$$\omega \in L \iff \exists C_{|\omega|} \forall x_1 M(\omega, x_1, C_{|\omega|}(\omega, x_1)) = 1.$$

It follows that $L \in \Sigma_2^p$. So as $\Pi_2^p \subseteq \Sigma_2^p$, we have that $\Sigma_2^p = \Pi_2^p$. So $\text{PH} = \Sigma_2^p$, as desired. \square

We conclude with Kannan's Theorem.

Theorem 227 (Kannan). For every polynomial $p(n)$, there exists a language $L \in \Sigma_2^p$ such that $L \notin \text{SIZE}(p(n))$.

The key idea is to prove this theorem for Σ_k^p , where k is some fixed integer greater than 2, and then apply the Karp-Lipton-Sipser Theorem to show that this language is actually in Σ_2^p .

Lemma 228. For every $k \in \mathbb{N}$, there exists a language $L \in \Sigma_4^p$ such that $L \notin \text{SIZE}(n^k)$.

Proof. The number of Boolean functions of the form $\varphi : \{0, 1\}^n \rightarrow \{0, 1\}$ that can be computed by circuits of size n^{k+1} is bounded above by $O(2^{n^{k+2}})$. So for n and k sufficiently large, there exists a set $S \subseteq \{0, 1\}^n$ that is not computed by any circuit of size n^{k+1} . On the other hand, by writing S in Disjunctive Normal Form, we may realize S with a circuit of size at most $(n+1) \cdot n^{2k} + n^{2k} \leq 2n^{2k+1}$.

We now define an ordering \prec on the set of circuits of size at most $2n^{2k+1}$, where $C \prec C'$ if the encoding of C is lexicographically smaller than the encoding of C' . Define L such that $L \cap \{0, 1\}^n$ is the subsets of $\{0, 1\}^n$ which satisfies the following:

- $L \cap \{0, 1\}^n$ is not accepted by a circuit of size n^{k+1} .
- $L \cap \{0, 1\}^n$ is accepted by the minimum (with respect to \prec) circuit \hat{C} of size $2n^{k+1}$.

We now claim that $L \in \Sigma_4^p$. We leave the proof as an exercise. □

We obtain the main theorem as a corollary. The proof is left as an exercise.

Corollary 229 (Kannan). For every $k \in \mathbb{N}$, there exists a language $L \in \Sigma_2^p$ such that $L \notin \text{SIZE}(n^k)$.

We next examine an extension of the Karp-Lipton-Sipser Theorem, due to Meyer.

Theorem 230 (Karp-Lipton-Meyer). Suppose that $\text{EXP} \subseteq \text{P/poly}$. Then $\text{EXP} = \Sigma_2^p$.

Proof. Let L be EXP-complete under polynomial-time reductions, and let M be the one-tape Turing Machine that accepts L and runs in time 2^{n^c} for some fixed c . Without loss of generality, suppose that M takes exactly 2^{n^c} steps. Suppose as well that, without loss of generality, $\Sigma = \{0, 1\}$. The key idea is to track the transcript of M . Define:

$$L_M := \{(x, i, j, z) : M(x) \text{ on step } i \text{ has a configuration whose } j\text{th bit is } z\}.$$

Here, $i \in [2^{n^c}]$, so i can be represented using n^c bits. Now any configuration of M uses at most 2^{n^c} bits. As j is an index of a string of length at most 2^{n^c} , we have that j can also be represented using n^c bits. So we can describe (x, i, j, z) using $\text{poly}(|x|)$ bits. Observe that $L_M \in \text{EXP} \subseteq \text{P/poly}$. So there exists a family $(C_n)_{n \in \mathbb{N}}$ of polynomial-sized circuits such that:

$$(x, i, j, z) \in L_M \iff C_{|x|}(x, i, j, z) = 1.$$

Note that as $i, j \in [2^{n^c}]$ and $z \in \{0, 1\}$, the size of the circuit depends only on $n := |x|$. By Lemma 226, there exists a polynomial-sized circuit C to compute the configuration of $M(x)$ at step i . Denote this configuration as $M_i(x)$, and let $(M_i(x))_j$ denote the j th bit of this configuration. Namely, C computes:

$$\langle (M_i(x))_1, \dots, (M_i(x))_{2^{n^c}} \rangle.$$

Now we have that $x \in L$ if and only if there is a polynomial-sized circuit C' , such that for all $i, j \in [2^{n^c}]$ and all $z \in \{0, 1\}$, $C(x, i, j, z) = 1 \iff (x, i, j, z) \in L_M$. It suffices to construct a polynomial-time verifier to check that $C(x, i, j, z) = 1 \iff (x, i, j, z) \in L_M$. As there is a polynomial-sized circuit C to compute the configuration of M at step i , we may in polynomial-time check the state and contents of the tape-head for M at step $i - 1$ and then check M 's transition function.

Thus, we have that $x \in L$ if and only if: there exists a circuit encoding $\langle C_{|x|} \rangle$ of polynomial-size, such that for all $i, j \in [2^{n^c}]$ and all $z \in \{0, 1\}$, we can verify in polynomial time that:

$$(x, i, j, z) \in L_M \iff C_{|x|}(x, i, j, z) = 1.$$

So $L \in \Sigma_2^p$, as desired. □

6.3.1 Exercises

(Recommended) Problem 84. Complete the proof of Lemma 228. Show that the language L constructed in the proof belongs to Σ_4^p .

(Recommended) Problem 85. We will prove Kannan's Theorem (Corollary 229): for every $k \in \mathbb{N}$, there exists a language $L \in \Sigma_2^p$ such that $L \notin \text{SIZE}(n^k)$.

- Show that if $\text{NP} \subseteq \text{P/poly}$, then the language L constructed in Lemma 228 satisfies the claim.
- Show that if $\text{NP} \not\subseteq \text{P/poly}$, then we may instead take $L = \text{SAT}$.

References

- [AA96] M. Agrawal and V. Arvind, *Geometric sets of low information content*, Theoretical Computer Science **158** (1996), no. 1, 193–219.
- [AB09] Sanjeev Arora and Boaz Barak, *Computational complexity: A modern approach*, 1st ed., Cambridge University Press, USA, 2009.
- [AKS02] Manindra Agrawal, Neeraj Kayal, and Nitin Saxena, *Primes is in p* , Ann. of Math **2** (2002), 781–793.
- [Bab15] László Babai, *Graph isomorphism in quasipolynomial time*, CoRR **abs/1512.03547** (2015).
- [BE99] Hans Ulrich Besche and Bettina Eick, *Construction of finite groups*, J. Symb. Comput. **27** (1999), no. 4, 387–404.
- [BEO02] Hans Ulrich Besche, Bettina Eick, and E.A. O’Brien, *A millennium project: Constructing small groups*, Intern. J. Alg. and Comput **12** (2002), 623–644.
- [BPG04] Paul Beame and Ethan Phelps-Goodman, *Lecture 10: Arthur-merlin games*, April 2004.
- [CH03] John J. Cannon and Derek F. Holt, *Automorphism group computation and isomorphism testing in finite groups*, J. Symb. Comput. **35** (2003), 241–267.
- [ELGO02] Bettina Eick, C. R. Leedham-Green, and E. A. O’Brien, *Constructing automorphism groups of p -groups*, Comm. Algebra **30** (2002), no. 5, 2271–2295. MR 1904637
- [Fil10] Yuval Filmus, *Smolensky’s lower bound*, September 2010.
- [Gro16] Joshua A. Grochow, *Np -hard sets are not sparse unless $p=np$: An exposition of a simple proof of mahaney’s theorem, with applications*, CoRR **abs/1610.05825** (2016).
- [GS86] S Goldwasser and M Sipser, *Private coins versus public coins in interactive proof systems*, Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC ’86, Association for Computing Machinery, 1986, p. 59–68.
- [Kat11] Jonathan Katz, *Notes on complexity theory- lecture 7*, Sep 2011.
- [Kha80] L. Khachiyan, *Polynomial algorithms in linear programming*, USSR Computational Mathematics and Mathematical Physics **20** (1980), 53–72.
- [KKL13] Swastick Koppertiy, John Kim, and Ben Lund, *Lecture 1: Course overview, circuits, and formulas*, 2013.
- [Lad75] Richard E. Ladner, *On the structure of polynomial time reducibility*, J. ACM **22** (1975), no. 1, 155–171.
- [Lev18] Michael Levet, *Fundamentals of computer science lecture notes*, 2018.
- [Lev20] Michael Levet, *Theory of computation- lecture notes*.
- [LGR16] François Le Gall and David J. Rosenbaum, *On the group and color isomorphism problems*, arXiv:1609.08253, 2016.
- [Mah82] Stephen R. Mahaney, *Sparse complete sets for np : Solution of a conjecture of berman and hartmanis*, Journal of Computer and System Sciences **25** (1982), no. 2, 130–143.
- [Mil78] Gary L. Miller, *On the $n \log n$ isomorphism technique (a preliminary report)*, Proceedings of the Tenth Annual ACM Symposium on Theory of Computing (New York, NY, USA), STOC ’78, Association for Computing Machinery, 1978, p. 51–58.
- [Mul] Wolfgang Mulzer, *The non-uniform hierarchy theorem*.
- [Ros12] Kenneth Rosen, *Discrete mathematics and its applications*, 7 ed., 2012.

- [Ros13] David J. Rosenbaum, *Bidirectional collision detection and faster deterministic isomorphism testing*, ArXiv **abs/1304.3935** (2013).
- [Sav97] John E. Savage, *Models of computation: Exploring the power of computing*, 1st ed., Addison-Wesley Longman Publishing Co., Inc., USA, 1997.
- [Sip96] Michael Sipser, *Introduction to the theory of computation*, 1st ed., International Thomson Publishing, 1996.
- [Vol99] Heribert Vollmer, *Introduction to circuit complexity: A uniform approach*, Springer-Verlag, Berlin, Heidelberg, 1999.
- [Wil19] James B. Wilson, *The threshold for subgroup profiles to agree is logarithmic*, Theory of Computing **15** (2019), no. 19, 1–25.