UNIVERSITY OF CALIFORNIA
Los Angeles

Communication and Complexity

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Computer Science

by

Pei Wu

2021

ABSTRACT OF THE DISSERTATION

Communication and Complexity

by

Pei Wu
Doctor of Philosophy in Computer Science
University of California, Los Angeles, 2021
Professor Alexander A. Sherstov, Chair

Communication is a universal process by which two or more individuals exchange information. A communication task that we study involves two (or more) parties, each given their own private input, trying to solve a function f or other computational task on their inputs together. For example, distributed data centers routinely check consistency with each other. An important feature of common communication phenomena is that it is interactive. The parties communicate alternatively and adaptively. This feature is absent in the classical information theory which focuses on one-way transmissions. However, the questions studied in the classical information theory remain valid and vital. For example, how do we minimize the communication cost for a certain communication task and how do we handle noise when it is present. From a historical point of view, theoretical computer scientists initiated the study of communication inspired mostly by complexity theory instead of information theory.

The latter source of inspiration motivates many other important questions. For example, how much more power a communication protocol gains given resources like randomness, nondeterminism, or quantum entanglement, etc.

In this dissertation, we discuss three concrete problems regarding the above questions. In particular, we study the following three problems.

(i) What is the maximum noise rate that can be tolerated in interactive communication? Specifically, we study the general noise model of arbitrary substitutions, deletions, and insertions. We settle this problem by giving a coding scheme that tolerates the maximum noise rate. A combinatorial ingredient of our scheme is that of *tree code*. We prove the existence of a tree code with strong distance properties.

(ii) What is the maximum communication complexity of constant-depth and polynomial-size circuits? We obtain strong communication lower bounds, ruling out the possibility of designing efficient generic communication protocols to solve this important class of problems. Our proof centers around the analytical measures *threshold degree* and *sign-rank*. The technique we use settles a 50-year-old problem in threshold degree that has applications to other areas of theoretical computer science including circuits complexity and learning theory.

(iii) How much power does a communication protocol gain taking advantage of quantum mechanics? We give a near-optimal separation between quantum and classical communication complexity, exhibiting functions that require only $O(\log n)$ bits of communication for quantum protocols but any classical protocol needs to essentially exchange the entire input. Our approach first studies the analogous problem in the *query model*, for which we are able to exhibit an optimal separation. These questions are broadly recognized as being central to understanding the phenomenon of quantum speedups.

The dissertation of Pei Wu is approved.

Raghu Meka

Amit Sahai

Rafail Ostrovsky

Alexander A. Sherstov, Committee Chair

University of California, Los Angeles

2021

Contents

List of Figures

Acknowledgments

During the journey of my career, a lot of thanks have already stacked on my head long before this dissertation was started.

Sasha Sherstov, my Ph.D. advisor, is the one who I would like to thank the most. In research, Sasha is just like a friend. We discuss whatever I am interested and I seek his wisdom whenever I feel in dark. Sasha has provided me with unlimited research ideas, and has the ability to see the essence and convert complicated problems into cleaner ones. I wanted to copy his skills badly. Sasha has been inspiring me not only by what he says but also by what he does. Sasha is extremely disciplined and proud about research (actually about everything he does). Besides many other things, these two philosophies influence me the most. People learn by imitating. Sasha shows us how a real researcher is like. Even outside research, Sasha is as supportive to his students. He seeks every opportunity to help me enjoy my life. The support I received from him is beyond words, and the fact that I know I have someone to rely on is the most blessed thing. I am very grateful to have Sasha as my advisor. Without Sasha, there will not be this dissertation.

Next, I would like to thank my committee members, they are Prof. Raghu Meka, Prof. Rafail Ostrovsky, and Prof. Amit Sahai. I would also like to thank Prof. Eli Gafni. They all have helped me in many different ways during my Ph.D. years.

I would like to thank many of my graduate colleagues at UCLA. They are, Prabhanjan Ananth, Yuan He, Aayush Jain, Ashutosh Kumar, Lun Liu, Paul Lou, Andrey

Storozhenko, Tianyi Zhang, Tao Zhou, Diyu Zhou. My graduate time here is more colorful because of you. Andrey, it is a quite enjoyable experience to work with you both as TAs for our CS 181 and for our research project. Ashutosh, we had some exciting and probably wild discussions. I hope one day it all come true. Yuan, Tianyi, and Diyu, clearly, we wasted a lot of time together. But it was fun. It is quite unfortunate to all of my UCLA colleagues that due to the pandemic, for most of us, we did not have much chance to meet and chat during the last and half years of my Ph.D. here. Nor did we have a proper farewell.

Before my years at UCLA, there were many people who greatly influenced me. First, I want to express my gratitude to my history teacher Wei Chen and computer science teacher Jun Yue in high school. Jun Yue taught me how to program in Pascal. That was how I learned about computer science properly and later decided to take it as my major at Nanjing University. At Nanjing University, I took lectures given by Prof. Yitong Yin. Besides many other things, Yitong's ideal of education left me with a strong impression. Finally, I would like to thank my Master's thesis advisor, Prof. Amit Chakrabarti, from Dartmouth College. Amit was the one who taught me communication complexity in the first place, which is the main focus of this dissertation.

My final thanks are to my parents. I have been away from home for so long. I wish I could have gone home more frequently and spent more time with them.

VITA

Education

2015–present UNIVERSITY OF CALIFORNIA, LOS ANGELES

Ph.D. candidate, Computer Science, Advisor: Alexander Sherstov

2013–2015 DARTMOUTH COLLEGE

M.S., Computer Science, Thesis Advisor: Amit Chakrabarti

2009–2013 NANJING UNIVERSITY, CHINA

Bachelor of Science, Computer Science and Technology

Publications

Optimal interactive coding for insertions, deletions, and substitutions

A. A. Sherstov, P. Wu

58th Annual Symposium on Foundations of Computer Science (FOCS 2017)

IEEE Transactions on Information Theory, **65**(10):5971–6000, 2019

Near-optimal lower bounds on the threshold degree and sign-rank of $\mathbf{AC}^0$

A. A. Sherstov, P. Wu

51st ACM Symposium on Theory of Computing (STOC 2019)

Invited to appear in SIAM Journal on Computing (special issue for STOC 2019)

An optimal separation of randomized and quantum query complexity

A. A. Sherstov, A. A. Storozhenko, P. Wu

53rd ACM Symposium on Theory of Computing (STOC 2021)

Speaking Engagements

Optimal interactive coding for insertions, deletions, and substitutions

- FOCS 2017, October 15-17, 2017 in Berkeley, California

Near-optimal lower bounds on the threshold degree and sign-rank of $\mathbf{AC}^0$

- STOC 2019, June 23-26, 2019 in Phoenix, Arizona
- Invited plenary talk, Feb, 2020 in UC Riverside, CA

An optimal separation of randomized and quantum query complexity

- QIP 2021 contributed talk (online), Feb 1-5, 2021
- Invited seminar talk (online), April, 2021, University of Waterloo, Canada
- STOC 2021 (online), June 21-25, 2021

Honors and Awards

6/2020	Outstanding Graduate Student Research Award (Computer Science Department, UCLA)
10/2020	Dissertation Year Fellowship (Graduate Division, UCLA)

Teaching and Service

Conference/journal reviewing: ICALP, STOC/FOCS, Algorithmica, SICOMP, TIT

Teaching assistant: CS 31 (Algorithms at Dartmouth College), CS 181 (Formal Language and Automata Theory at UCLA)

CHAPTER 1

Introduction

Communication is a universal process by which two or more individuals exchange information. The rapid evolution in modern communication technology from mobile devices to the Internet has dramatically changed the world. It will continuously bring social values. No need to be loquacious, communication is important on its own. For a computer scientist, communication is also interesting as it is closely related to computation. In any computational device, there are components in the device that communicate in disguise with each other during all sorts of computation. Communication is not simply a product of specific algorithms but is inherent in many situations. Due to this reason, the study of communication provides a way to understand computation. Driven by its practical and theoretical importance, the study of communication has evolved into a large body of research.

In this dissertation, we study the communication phenomenon from the theoretical side. We focus on the abstract communication model. The basic two-party model, proposed by Yao [**139**], features two geographically separated parties, Alice and Bob. Alice and Bob have private inputs x and y, respectively, and need to communicate back and forth to accomplish some communication task, say to compute a given function $f(x, y)$. We study the communication between Alice and Bob and ignore any computation taken by each individual. The *communication channel* between Alice and Bob may be classical, may be quantum, and may suffer from adversarial noise. The protocol may need to output a correct answer all the time, may need to output a correct answer with *bounded error*, or with *unbounded error*. We will discuss all these aspects in detail in later chapters.

There are two perspectives we take on the study of communication. The first one is information theoretical. Classical information theory studies one-way transmission of information as opposed to interactive communication. Two of the fundamental problems studied by classical information theory are: what is the minimum communication cost to transmit a message? how does one handle noise if it is present? It is natural to generalize the theory to an interactive setting. Under this point of view, there are two analogous questions that we can ask:

Q1. For a given function $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$, what is the minimum communication cost for Alice and Bob to solve f?

Q2. Suppose we already have a protocol π designed for the noiseless channel, how to encode π such that the encoded protocol can tolerate noise?

There are other fundamental questions, for example, how do we define the information content of a protocol, and do we have a general compression theorem that compresses a protocol to its information content. These questions are out of the scope of this dissertation.

The second perspective is complexity theoretical. A goal of computational complexity theory is to understand the limitations of all sorts of computing devices with provided resources. Communication is a resource. The study of *communication complexity* is to prove lower bounds of communication cost for certain communication problems. Therefore, communication complexity answers the first question in the previous paragraph from the lower bound side. This new perspective also has a far-reaching influence on computer science. The reason is already mentioned—computation and communication are intertwined. Major breakthroughs in various sub-areas in computer science are made possible via the lens of communication complexity. In fact, communication complexity has applications in almost all areas of theoretical computer science, including circuit complexity, streaming algorithms, learning theory, etc. Another goal of complexity theory is to understand the relative power of different complexity classes. An example is the well-known P v.s. NP problem. Inspired by this goal from complexity theory, we can ask similar questions in the communication context, like

Q3. How much more do power communication protocols gain given nondeterminism, randomness, or quantum entanglement?

An intriguing aspect is that many questions can be answered in the communication model but the analogous questions are beyond reach in any near future for Turing machines.

1.1. Our contributions

In this section, we put the three questions mentioned above in concrete settings. We will briefly discuss our discoveries regarding these questions.

(i) What is the maximum noise rate that can be tolerated in interactive communication? Specifically, we study the general noise model of arbitrary substitutions, deletions, and insertions. We settle this problem by giving a coding scheme that tolerates the maximum noise rate. A key combinatorial ingredient of our scheme is that of tree code. We prove the existence of a tree code with strong distance properties.

(ii) What is the maximum communication complexity of constant-depth and polynomial-size circuits? Our focus is the unbounded-error regime. We obtain an essentially optimal lower bound. Our proof centers around the analytical measures threshold degree and sign-rank. The technique we use settles a 50-year-old problem in threshold degree that has applications to other areas of theoretical computer science including circuit complexity and learning theory.

(iii) How much power does a communication protocol gain taking advantage of quantum mechanics? We give a near-optimal separation between quantum and classical communication complexity, exhibiting functions that require only $O(\log n)$ bits of communication for quantum protocols but any classical communication protocol needs to essentially exchange the entire input. Our approach first studies the analogous problem in the *query model*, for which we are able to exhibit and optimal separation. These questions are broadly recognized as being central to understanding the phenomenon of quantum speedups.

1.1.1. Interactive coding. Noise is omnipresent in communication. In the classical setting of one-way communication, the study of information transmission under noise forms a large part of classical information theory. In pioneering work, Schulman [**111**] considered noise in the setting of interactive communication. This area of research, called interactive coding, is a fascinating and highly active discipline at the crossroads of information theory and communication complexity. More concretely, consider the following scenario. Alice and Bob would like to execute a communication protocol π defined for a noiseless environment. However, the communication channel is controlled by an adversary who can substitute any fraction ρ of symbols transmitted through the channel. The question is, can Alice and Bob use some interactive analogue of error-correcting codes to ensure that they are both able to recover, from their noisy communication, the transcript that π would have been produced without noise? Braverman and Rao [**26**] gave an affirmative answer for any constant $\rho \in [0, 1/4)$.

A far-reaching generalization of this model, proposed by Braverman et al. [**25**], allows *arbitrary* corruptions: insertions, deletions, and substitutions. Besides being a natural and interesting generalization that has been studied extensively in coding theory [**87, 112, 70**], this corruption model has also found application in the standard corruption model (substitution errors only) [**21**]. For any constant $\varepsilon > 0$, the authors of [**25**] showed how to faithfully simulate any protocol in this generalized model with corruption rate up to $\frac{1}{18} - \varepsilon$, using a constant-size alphabet and a constant-factor overhead in communication.

Braverman et al. [**25**]'s work left open the following natural and fundamental question: what is the maximum corruption rate that can be tolerated in this generalized model of substitutions, insertions, and deletions? We gave a complete and somewhat surprising answer to this question in [**128**]. We showed that for any $\varepsilon > 0$, there is an interactive coding scheme that uses a constant-size alphabet and achieves the noise tolerance rate of $\frac{1}{4} - \varepsilon$, at the expense of a constant-factor overhead in communication complexity compared to π. This rate is easily seen to be optimal, even in the presence of substitution errors alone.

1.1.2. Unbounded-error communication. Like many other computational models, a protocol is allowed to err when taking advantage of randomness. For the protocol to be meaningful computing a Boolean function, however, its error probability should be strictly smaller than 1/2. If we only enforce such a criterion of success on the communication protocol, then its minimum communication cost is called the communication complexity with *unbounded error*, unbounded in the sense that the error probability can be arbitrarily close but not equal to 1/2. This model is first defined and studied by Paturi and Simon [**101**]. As one would expect, the relaxed success criterion grants the protocol more power, and therefore it is harder to prove strong lower bounds in this model. In fact, the communication complexity with unbounded error is one of the strongest communication models (e.g., stronger than randomized and quantum communication model) for which we have tools to prove explicit lower bounds. For this reason, the study of communication protocols with unbounded error stands at the frontier of current research. *Constant depth circuits* ($\mathbf{AC}^0$) is a central complexity class in theoretical computer science. Many of the greatest achievements in theoretical computer science address questions regarding $\mathbf{AC}^0$. In the pioneering work of Babai et al. [**11**], the authors asked the question: are there functions in $\mathbf{AC}^0$ that do not admit any efficient communication protocols with unbounded error.

We prove that there are functions computable by constant depth circuits such that Alice (or Bob) has to essentially send her entire input for them to achieve a success probability barely larger than 1/2. This is the strongest communication lower bounds for $\mathbf{AC}^0$. In the complexity theoretical language, our result optimally separates the *polynomial hierarchy* PH and communication complexity class with unbounded error UPP. The precise definitions of these complexity classes can be found in [**11**], and are analogues of the corresponding complexity classes for Turing machines.

Our approach centers around the analysis of the sign-representation of Boolean functions. Representations of Boolean functions by real polynomials are of great importance in many different contexts, from communication complexity and quantum computing to machine learning theory. For example, the notion of *approximate degree* has played an essential role in quantum query complexity for decades. *Threshold degree* has an even broader range of applications, including various models of computational learning. The notion of threshold degree originates in the pioneering work of Minsky and Papert [**90**] and is defined, for a Boolean function $f : \{0,1\}^n \to \{-1,1\}$, as the minimum degree of a real polynomial p that represents f in sign: $(-1)^{f(x)} = \operatorname{sgn} p(x)$ for all x. In particular, the threshold degree of polynomial-size constant-depth circuits has been the focus of 50 years of work. We were able to essentially settle this long-standing problem in [**129**]. More specifically, we proved that for any $\varepsilon > 0$, there is an $\mathbf{AC}^0$ circuit with threshold degree $\Omega(n^{1-\varepsilon})$. This lower bound essentially matches the trivial upper bound of $O(n)$ and is a polynomial improvement on the best previous lower bound, $\Omega(\sqrt{n})$. We also proved a much stronger result that applies not only to threshold degree but also to *sign-rank*—a vastly more general notion than threshold degree. A remarkable fact due to Paturi and Simon is that the sign-rank and the communication complexity with unbounded error are essentially the same mathematical notion [**101**].

Among many other things, our results have important applications in learning theory. In particular, our results rule out the possibility of the distribution-free PAC learning of $\mathbf{AC}^0$ based on the powerful dimension complexity paradigm. This framework captures nearly all known algorithmic results for distribution-free PAC learning.

1.1.3. Quantum versus classical communication. One central task in quantum computing is to answer how much more powerful quantum computers can be than classical computers. Of particular prominence in this line of research are results demonstrating the superiority of quantum algorithms over their classical counterparts. For example, the celebrated Shor's factoring algorithm shows how to factor a number in polynomial time by quantum computers, for which no classical polynomial-time algorithm is known despite decades of research. In the context of communication, it is trivial that there is a function whose deterministic communication complexity is linear but quantum communication complexity is constant, e.g., the equality problem. So the more interesting question is: what is the largest separation between the randomized and quantum communication complexity. We are particularly interested in the *bounded-error* regime, where the protocol is allowed to err with a small constant probability. The first exponential separation was due to Raz [**103**], followed by a sequence of improvement [**107, 57, 135**].

We prove in our work that for any constant $\varepsilon > 0$, there is a *partial function* f, such that the quantum communication complexity of f is $O(\log n)$ but its randomized communication complexity is $\Omega(n^{1-\varepsilon})$ [**127**]. This separation is essentially optimal and a polynomial improvement on previous work. On the other hand, if we are only interested in *total functions*, we prove a cubic separation, again improving prior works by a polynomial factor. Notice how small this gap is compared to the case of partial functions. It is a major open problem to determine if for total functions quantum and randomized communication complexity are in fact polynomially related.

We obtain the above separation results by first proving an analogous result in the query model, and then apply a very general framework that lifts our hardness results in the query model to the communication model. Quantum query complexity has been studied even more extensively and can be justly considered to be among the biggest achievements of quantum computing to date. In groundbreaking work, Simon [**131**] exhibited a partial Boolean function whose bounded-error quantum query complexity is exponentially smaller than its randomized (i.e., classical) query complexity. This raises the question: what is the largest possible separation between quantum and randomized query complexity? This question was first explicitly stated in 2002 by Buhrman et al. [**30**], and has since been popularized by Aaronson and Ambainis [**2**]. We settled this 18-year-old problem completely in [**127**]. Specifically, we proved that for any constant k, there is a partial function f with quantum query complexity at most k and randomized query complexity $\tilde{\Omega}(n^{1-1/2k})$. This gives an $O(1)$ versus $\Omega(n^{1-\varepsilon})$ separation for any $\varepsilon > 0$, which is a polynomial improvement on the best previous separation of $O(1)$ versus $\Omega(n^{2/3-\varepsilon})$ due to Tal [**135**]. Moreover, our separation is optimal due to Aaronson and Ambainis' result [**2**] that any k-query quantum algorithm can be simulated by $O(n^{1-1/2k})$ randomized classical queries, for an arbitrary constant k. By the well-known framework of "cheatsheets" due to Aaronson et al. [**3**], our result also implies a cubic separation between quantum and randomized query complexity for *total* functions. This separation is the largest known and has been conjectured to be tight by other researchers [**4**].

As a technical centerpiece of our work [**127**], we prove a tight bound on the ℓ_1 norm of any given level of the Fourier spectrum of decision trees. This bound on Fourier weight settles a conjecture of Tal [**135**] and is of substantial interest in its own right, considering the central role of the Fourier spectrum in many recent breakthroughs in the area [**45, 46, 104**].

1.2. Organization

In Chapter 2, we set our notations and provide a basic mathematical background for the remaining chapters. In Chapter 3, we define various communication models in detail. From deterministic communication, randomized communication with bounded error or unbounded error, to the more advanced quantum communication and multiparty communication. We also discuss a few general techniques that we will use to show strong communication lower bounds for various models. In Chapter 4, we discuss our contribution to communication against the adversarial noise of insertions, deletion, and substitutions based on the paper [**128**]. This is a joint work with Sherstov. In Chapter 5, we present our work on the unbounded-error communication of $\mathbf{AC}^0$. This part is based on the paper [**129**], another joint work with Sherstov. In Chapter 6, we discuss our near-optimal result on the separation of quantum and classical communications, exhibiting the superiority of quantum computation in the context of interactive communication as well as of the query model. This is based on the paper [**127**], a joint work with Sherstov and Storozhenko. In the final chapter, we conclude our dissertation and point out some problems for future research.

CHAPTER 2

Notations and preliminaries

In this chapter, we set the notations, and provide a minimum background on the technical tools we will use in this dissertation. Most of the notations are standard, but some can be unfamiliar.

2.1. Numbers, sets and functions

We adopt the standard notation $\mathbb{N} = \{0, 1, 2, 3, \ldots\}$, $\mathbb{Z}^+ = \{1, 2, 3, \ldots\}$, $\mathbb{R}$, $\mathbb{R}^+$ and $\mathbb{C}$ for the sets of natural numbers, positive integers, real numbers, positive real numbers and complex numbers, respectively. We adopt the extended real number system $\mathbb{R} \cup \{-\infty, \infty\}$ in all calculations. In particular, we have $a/0 = \infty$ for any positive number $a \in \mathbb{R}$. To simplify our notation, we further adopt the convention that

$$\frac{0}{0} = 0.$$

For any $a, b \in \mathbb{R}$, we adopt the standard notation of open/closed intervals, i.e.,

$$\begin{aligned}
(a, b) &= \{x \in \mathbb{R} : a < x < b\}, \\
[a, b] &= \{x \in \mathbb{R} : a \leq x \leq b\}, \\
[a, b) &= \{x \in \mathbb{R} : a \leq x < b\}, \\
(a, b] &= \{x \in \mathbb{R} : a < x \leq b\}.
\end{aligned}$$

Occasionally, we will also encounter finite fields, e.g., $\mathbb{F}_2 = \{0, 1\}$. Number 0 and 1 are commonly used to represent the "true" or "false" value of a Boolean variable. Another common representation is $\{-1,1\}$, where -1 encodes "true" and 1 encodes "false." Each representation can be convenient in different scenarios. The default representation in this dissertation is $\{0,1\}$, but we will also use the $\{-1,1\}$ representation. We use the standard definition of the sign function:

$$\operatorname{sgn} x = \begin{cases} -1 & \text{if } x < 0, \\ 0 & \text{if } x = 0, \\ 1 & \text{if } x > 0. \end{cases}$$

As usual, the *cardinality* of a set A, denoted $|A|$, is the number of elements in A. The complement of a set A is denoted $\overline{A}$. For arbitrary sets A and B, we define the *cardinality of A relative to B* by $|A|_B = |A \cap B|$. For a set A and a sequence s of elements, we let $A \cup s$ denote the set of elements that occur in either A or s. We define $A \cap s$ analogously. For nonempty sets $A, B \subseteq \mathbb{R}$, we write $A < B$ to mean that $a < b$ for all $a \in A$, $b \in B$. It is clear that this relation is a partial order on nonempty subsets of $\mathbb{R}$. The power set of a set S, denoted by $\mathcal{P}(S)$, is the set of all subsets of S. For a set S and a nonnegative integer k, we let $\binom{S}{k}$ denote the family of subsets of S that have cardinality exactly k, and $\binom{S}{<k}$ the family of subsets of S that have cardinality less than k. $\binom{S}{\leq k}$, $\binom{S}{>k}$, $\binom{S}{\geq k}$ are defined analogously We further define

$$\mathcal{P}_{n,k} = \binom{\{1,2,\ldots,n\}}{k} = \{S \subseteq \{1,2,\ldots,n\} : |S| = k\}.$$

The following well-known bound [**73**, Proposition 1.4] is used in our proofs without further mention:

$$\left(\frac{n}{k}\right)^k \leq \binom{n}{k} \leq \left(\frac{en}{k}\right)^k, \qquad k = 1,2,\ldots,n, \tag{2.1.1}$$

where $e = 2.7182\ldots$ denotes Euler's number.

For a logical condition C, we use the Iverson bracket:

$$\mathbf{I}[C] = \begin{cases} 1 & \text{if } C \text{ holds,} \\ 0 & \text{otherwise.} \end{cases}$$

For a finite set X, we let $\mathbb{R}^X$ denote the family of real-valued functions on X. The *support* of a function $f \in \mathbb{R}^X$ is denoted $\operatorname{supp} f = \{x \in X : f(x) \neq 0\}$. For real-valued functions with finite support, we adopt the usual norms and inner product:

$$\begin{aligned}
\|f\|_\infty &= \max_{x \in \operatorname{supp} f} |f(x)|, \\
\|f\|_1 &= \sum_{x \in \operatorname{supp} f} |f(x)|, \\
\|f\|_2 &= \left(\sum_{x \in \operatorname{supp} f} |f(x)|^2 \right)^{1/2}, \\
\langle f, g \rangle &= \sum_{x \in \operatorname{supp} f \cap \operatorname{supp} g} f(x)g(x).
\end{aligned}$$

We adopt the convention that $\|f\| = \|f\|_2$. This covers as a special case functions on finite sets. For $f, g \in \mathbb{R}^X$, we let $f \cdot g \in \mathbb{R}^X$ denote the point-wise product of f and g, with $(f \cdot g)(x) = f(x)g(x)$. For a real-valued function $f\colon X \to \mathbb{R}$, recall that $\arg\min_{x \in X} f(x)$ denotes the set of points where f attains its minimum value. Analogously, $\arg\max_{x \in X} f(x)$ denotes the set of points where f attains its maximum value.

The functions $\ln x$ and $\log x$ stand for the natural logarithm of x and the logarithm of x to base 2, respectively. To avoid excessive use of parentheses, we follow the notational convention that $\ln a_1 a_2 \ldots a_k = \ln(a_1 a_2 \ldots a_k)$ for any factors $a_1, a_2, \ldots, a_k$. The binary entropy function $H\colon [0,1] \to [0,1]$ is given by

$$H(x) = x \log \frac{1}{x} + (1 - x) \log \frac{1}{1 - x}.$$

Basic calculus reveals that

$$H(x) \leq 1 - \frac{2}{\ln 2}\left(x - \frac{1}{2}\right)^2. \tag{2.1.2}$$

2.2. Strings

In this dissertation, an *alphabet* Σ is any nonempty finite set of symbols other than the asterisk $*$, which we treat as a reserved symbol. Recall that Σ^* stands for the set of all strings over Σ. We denote the empty string as usual by ε. For an alphabet Σ and a natural number n, we let Σ^n denote the set of all string over Σ of length n, and $\Sigma^{\leq n}$ denote the set of strings of length up to n, so that $\Sigma^{\leq n} = \{\varepsilon\} \cup \Sigma \cup \Sigma^2 \cup \cdots \cup \Sigma^n$. For any alphabet Σ, we let $\prec$ denote the standard partial order on Σ^* whereby $u \prec v$ if and only if $uw = v$ for a nonempty string w. The derived relations $\succ, \preceq, \succeq$ are defined as usual by

$$\begin{aligned} u \succ v &\quad\Leftrightarrow\quad v \prec u, \\ u \succeq v &\quad\Leftrightarrow\quad v \prec u \text{ or } v = u, \\ u \preceq v &\quad\Leftrightarrow\quad u \prec v \text{ or } v = u. \end{aligned}$$

A *prefix* of v is any string u with $u \preceq v$. A *suffix* of v is any string u such that $v = wu$ for some string w. A prefix or suffix of v is called *proper* if it is not equal to v. A *subsequence* of v is v itself or any string that can be obtained from v by deleting one or more symbols.

For a string v over a given alphabet, we let $|v|$ denote the length of v. For a set S, we let $v|_S$ denote the substring of v indexed by the elements of S. In other words, $v|_S = v_{i_1} v_{i_2} \cdots v_{i_{|S|}}$ where $i_1 < i_2 < \cdots < i_{|S|}$ are the elements of S. For a number $\iota \in [0, \infty]$ in the extended real number system, we let $v_{<\iota}$ denote the substring of v obtained by keeping the symbols at indices less than ι. As special cases, we have

$v_{<1} = \varepsilon$ and $v_{<\infty} = v$. The substrings $v_{\leq\iota}$, $v_{>\iota}$, and $v_{\geq\iota}$ are defined analogously. In any of these four definitions, an index range that is empty produces the empty string ε.

2.3. Vectors and matrices

The abstract vector space is a commutative group $\mathcal{A}$ over a field $\mathbb{F}$ satisfying laws of distributivity. We will work only with the Euclidean vector space and its subspaces. The n-dimensional Euclidean space is the vector space $\mathbb{R}^n$ equipped with inner product $\langle\cdot,\cdot\rangle$. Its standard basis is the set $\{e_1, e_2, \ldots, e_n\}$, where e_i is the vector whose coordinates are all 0 except the ith coordinate which is 1. Analogous to functions, we adopt the familiar norms for vectors $x \in \mathbb{R}^n$ in Euclidean space:

$$\|x\|_\infty = \max_{i=1,\ldots,n} |x_i|,$$

$$\|x\|_1 = \sum_{i=1}^{n} |x_i|,$$

$$\|x\|_2 = \left(\sum_{i=1}^{n} |x_i|^2\right)^{1/2}.$$

To avoid notational clutter we use $|x|$ interchangeably with $\|x\|_1$[1]. We refer to $|x| = \|x\|_1$ as the *weight* of x. We also often omit the subscript in the ℓ_2-norm, i.e., $\|x\| = \|x\|_2$. For vectors $x \in \mathbb{R}^n, y \in \mathbb{R}^m$, define partial orders $<$, i.e., $x < y$ means $x_i < y_i$ on all coordinates i. Analogously, we define $>, \leq, \geq$. Finally $x \gneq y$ means $x \geq y$ and $x \neq y$. The tensor product $x \otimes y$ denotes the vector $(\ldots, x_i y_j, \ldots) \in \mathbb{R}^{nm}$. The notation $x^{\otimes n}$ is the abbreviation for

$$\underbrace{x \otimes x \otimes \cdots \otimes x}_{n}.$$

[1] For a function f, $|f|$ is the function whose evaluation at x is $|f(x)|$.

Analogously, for vector space $\mathcal{A}$ with basis $\{\alpha_1, \alpha_2, \ldots, \alpha_n\}$ and vector space $\mathcal{B}$ with basis $\{\beta_1, \beta_2, \ldots, \beta_m\}$, the tensor product $\mathcal{A} \otimes \mathcal{B}$ is the vector space of dimension nm, whose basis is the set

$$\{(\alpha_i, \beta_j) : i = 1, 2, \ldots, n, j = 1, 2, \ldots, m\}.$$

$\mathcal{A}^{\otimes n}$ is defined in an analogous way for vector space $\mathcal{A}$.

For any vector space $\mathcal{A}$ over $\mathbb{R}$ and any subset S of $\mathcal{A}$, $\operatorname{span} S$ is the *subspace* spanned by vectors in S. In particular,

$$\operatorname{span} S = \left\{ \sum_{i=1}^{m} a_i x_i : a_i \in \mathbb{R}, x_i \in S, m \in \mathbb{N} \right\}.$$

In addition, we define the convex hull $\operatorname{conv} S$ and conical hull $\operatorname{cone} S$,

$$\begin{aligned} \operatorname{conv} S &= \left\{ \sum_{i=1}^{m} a_i x_i : a_i \in \mathbb{R}^+, \sum_{i=1}^{m} a_i = 1, x_i \in S, m \in \mathbb{N} \right\}, \\ \operatorname{cone} S &= \left\{ \sum_{i=1}^{m} a_i x_i : a_i \in \mathbb{R}^+, x_i \in S, m \in \mathbb{N} \right\}. \end{aligned}$$

In the context of quantum computing, the bra-ket notation is adopted. The ket $|\phi\rangle$ denotes a column vector, while the bra $\langle\psi|$ denotes a row vector. This notation can be memorized by inner product of two vectors $\langle\psi|\phi\rangle$ which can be interpreted as the matrix multiplication of a row vector and a column vector. For two vectors $|\phi\rangle, |\psi\rangle$, there are various notations used to denote their tensor product, e.g., $|\phi, \psi\rangle, |\phi\psi\rangle, |\phi\rangle|\psi\rangle$.

The symbol $\mathbb{R}^{n\times m}$ refers to the family of all n by m matrices with real entries. For any finite sets X, Y, we also use $\mathbb{R}^{X\times Y}$ denote the family of real matrices whose rows and columns are indexed by elements from X and Y, respectively. We specify matrices by their generic entry, e.g., $M = [f(x, y)]_{x\in X, y\in Y}$. We denote the rank of any matrix

$M \in \mathbb{R}^{n\times m}$ by $\operatorname{rk} M$. Recall that the spectral norm of M is given by

$$\|M\| = \max_{x\in\mathbb{R}^m : \|x\|=1} \|Mx\|.$$

2.4. Boolean functions and circuits

We view Boolean functions as mappings $X \to \{0,1\}$ for some finite set X. More generally, we consider *partial* Boolean functions $f\colon X \to \{0,1,*\}$, with the output value $*$ used for don't-care inputs. The negation of a Boolean function f is denoted as usual by $\neg f = 1 - f$. The familiar functions $\mathrm{OR}_n, \mathrm{AND}_n, \mathrm{XOR}_n, \mathrm{MAJ}_n : \{0,1\}^n \to \{0,1\}$ are given by

$$\begin{aligned}
\mathrm{OR}_n(x) = 1 \qquad &\Leftrightarrow \qquad \exists\, i, x_i = 1,\\
\mathrm{AND}_n(x) = 1 \qquad &\Leftrightarrow \qquad \forall\, i, x_i = 1,\\
\mathrm{XOR}_n(x) = 1 \qquad &\Leftrightarrow \qquad \sum_{i=1}^n x_i \text{ is odd},\\
\mathrm{MAJ}_n(x) = 1 \qquad &\Leftrightarrow \qquad \sum_{i=1}^n x_i > n/2.
\end{aligned}$$

We also refer to XOR as the parity function, as it computes the parity of input's weight. We abbreviate $\mathrm{NOR}_n = \neg\mathrm{OR}_n$.

We adopt the standard notation for function composition, with $f \circ g$ defined by $(f \circ g)(x) = f(g(x))$. In addition, we use the $\circ$ operator to denote the *component-wise* composition of Boolean functions. Formally, the component-wise composition of $f\colon \{0,1\}^n \to \{0,1\}$ and $g\colon X \to \{0,1\}$ is the function $f \circ g\colon X^n \to \{0,1\}$ given by $(f \circ g)(x_1, x_2, \ldots, x_n) = f(g(x_1), g(x_2), \ldots, g(x_n))$. Component-wise composition is consistent with standard composition, which in the context of Boolean functions is only defined for $n = 1$. Thus, the meaning of $f \circ g$ is determined by the range of g and is never in doubt. Component-wise composition generalizes in the natural manner to partial Boolean functions $f\colon \{0,1\}^n \to \{0,1,*\}$ and $g\colon X \to \{0,1,*\}$, as follows:

$$(f \circ g)(x_1,\ldots,x_n) = \begin{cases} f(g(x_1),\ldots,g(x_n)) & \text{if } x_1,\ldots,x_n \in g^{-1}(0 \cup 1), \\ * & \text{otherwise.} \end{cases}$$

Compositions $f_1 \circ f_2 \circ \cdots \circ f_k$ of three or more functions, where each instance of the $\circ$ operator can be standard or component-wise, are well-defined by associativity and do not require parenthesization.

For Boolean strings $x, y \in \{0,1\}^n$, we let $x \oplus y$ denote their bitwise XOR. The strings $x \wedge y$ and $x \vee y$ are defined analogously, with the binary connective applied bitwise. A *Boolean circuit* C in variables $x_1, x_2, \ldots, x_n$ is a circuit with inputs $x_1, \neg x_1, x_2, \neg x_2, \ldots, x_n, \neg x_n$ and gates $\wedge$ and $\vee$. The circuit C is *monotone* if it does not use any of the negated inputs $\neg x_1, \neg x_2, \ldots, \neg x_n$. The *fan-in* of C is the maximum in-degree of any $\wedge$ or $\vee$ gate. Unless stated otherwise, we place no restrictions on the gate fan-in. The *size* of C is the number of $\wedge$ and $\vee$ gates. The *depth* of C is the maximum number of $\wedge$ and $\vee$ gates on any path from an input to the output gate. With this convention, the circuit that computes $(x_1, x_2, \ldots, x_n) \mapsto x_1$ has depth 0. The circuit class $\mathbf{AC}^0$ consists of function families $\{f_n\}_{n=1}^{\infty}$ such that each $f_n\colon \{0,1\}^n \to \{0,1\}$ is computed a Boolean circuit of size at most cn^c and depth at most c, for some constant $c \geq 1$ and all n. We specify small-depth layered circuits by indicating the type of gate used in each layer. For example, an *AND-OR-AND circuit* is a depth-3 circuit with the top and bottom layers composed of $\wedge$ gates, and middle layer composed of $\vee$ gates. A *Boolean formula* is a Boolean circuit in which every gate has fan-out 1. Common examples of Boolean formulas are DNF and CNF formulas.

2.5. Fourier transform

In this section, we discuss Fourier transform of real-valued functions whose domain is the finite Abelian group of the Boolean cubes, e.g., $f : \{0,1\}^n \to \mathbb{R}$. We use the $\{0,1\}^n$ representation at first for definitions. However, all the discussions apply to the $\{-1,1\}^n$ representation in a straightforward manner.

Consider the real vector space of functions $\{0,1\}^n \to \mathbb{R}$. For $S \subseteq \{1,2,\ldots,n\}$, define $\chi_S\colon \{0,1\}^n \to \{-1,1\}$ by $\chi_S(x) = (-1)^{\sum_{i\in S} x_i}$. Then

$$\langle \chi_S, \chi_T \rangle = \begin{cases} 2^n & \text{if } S = T, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, $\{\chi_S\}_{S\subseteq\{1,2,\ldots,n\}}$ is an orthogonal basis for the vector space in question. In particular, every function $\phi\colon \{0,1\}^n \to \mathbb{R}$ has a unique representation of the form

$$\phi = \sum_{S\subseteq\{1,2,\ldots,n\}} \hat{\phi}(S)\chi_S$$

for some reals $\hat{\phi}(S)$, where by orthogonality $\hat{\phi}(S) = 2^{-n}\langle \phi, \chi_S \rangle$. The reals $\hat{\phi}(S)$ are called the *Fourier coefficients of* ϕ, and the mapping $\phi \mapsto \hat{\phi}$ is the *Fourier transform of* ϕ. The *order* of a Fourier coefficient $\hat{\phi}(S)$ is the cardinality $|S|$. The *degree* of ϕ is the quantity $\max\{|S| : \hat{\phi}(S) \neq 0\}$. From basic linear algebra, one can derive Plancherel's theorem

$$\langle f, g \rangle = 2^n \sum_{S\subseteq\{1,2,\ldots,n\}} \hat{f}(S)\hat{g}(S).$$

A simple and useful corollary (Parseval's identity) is that $\sum_S \hat{f}^2(S) = 1$ if $f(x)$ takes values only in $\{-1,1\}$, i.e., $f : \{0,1\}^n \to \{-1,1\}$. The following fact is also immediate from the definition of $\hat{\phi}(S)$.

Proposition 2.1. *Let $\phi\colon \{0,1\}^n \to \mathbb{R}$ be given. Then*

$$\max_{S\subseteq\{1,2,\ldots,n\}} |\hat{\phi}(S)| \leq 2^{-n}\|\phi\|_1.$$

For $k = 0,1,2,\ldots,n$, we introduce the linear operator $L_k\colon \mathbb{R}^{\{0,1\}^n} \to \mathbb{R}^{\{0,1\}^n}$ that sends a function $\phi\colon \{0,1\}^n \to \mathbb{R}$ to the function $L_k\,\phi\colon \{0,1\}^n \to \mathbb{R}$ given by

$$(L_k\,\phi)(x) = \sum_{S\in\mathcal{P}_{n,k}} \hat{\phi}(S)\chi_S(x).$$

We refer to $L_k\,\phi$ as the *degree-k homogeneous part of ϕ.*

Note that when using the $\{-1,1\}$ representation, by the above discussion every function $\phi : \{-1,1\}^n \to \mathbb{R}$ has a unique representation as a multilinear polynomial

$$\phi(x) = \sum_{S\subseteq\{1,2,\ldots,n\}} \hat{\phi}(S)\prod_{i\in S} x_i, \tag{2.5.1}$$

where the real numbers $\hat{\phi}(S)$ are the Fourier coefficients of f.

For any polynomial $p \in \mathbb{R}[x_1,x_2,\ldots,x_n]$, we let $|||p|||$ denote the sum of the absolute values of the coefficients of p. One easily verifies the well-known fact that $|||\cdot|||$ is a norm on the polynomial ring $\mathbb{R}[x_1,x_2,\ldots,x_n]$. We identify a function $\phi\colon \{-1,1\}^n \to \mathbb{R}$ with its unique representation (2.5.1) as a multilinear polynomial, to the effect that

$$|||\phi||| = \sum_{S\subseteq\{1,2,\ldots,n\}} |\hat{\phi}(S)|$$

is the sum of the absolute values of the Fourier coefficients of ϕ. We will abuse the notation $|||\phi|||$ as the sum of absolute values of the Fourier coefficients of ϕ under $\{0,1\}$ representation as well.

PROPOSITION 2.2. *For any functions $\phi, \psi : \{0,1\}^n \to \mathbb{R}$ and reals a, b,*

$$\|\!|a\phi + b\psi\|\!| \leq |a| \, \|\!|\phi\|\!| + |b| \, \|\!|\psi\|\!|.$$

Proof. We have

$$\begin{aligned}\|\!|a\phi + b\psi\|\!| &= \sum_{S\subseteq\{1,2,\dots,n\}} |a\hat{\phi}(S) + b\hat{\psi}(S)| \\ &\leq |a| \sum_{S\subseteq\{1,2,\dots,n\}} |\hat{\phi}(S)| + |b| \sum_{S\subseteq\{1,2,\dots,n\}} |\hat{\psi}(S)| \\ &= |a| \, \|\!|\phi\|\!| + |b| \, \|\!|\psi\|\!|,\end{aligned}$$

where the first step uses the linearity of the Fourier transform. □

We also note the following submultiplicative property.

PROPOSITION 2.3. *For any functions $\phi, \psi\colon \{0,1\}^n \to \mathbb{R}$,*

$$\|\!|\phi \cdot \psi\|\!| \leq \|\!|\phi\|\!| \, \|\!|\psi\|\!|.$$

Proof. We have

$$\begin{aligned}\phi \cdot \psi &= \left(\sum_{S\subseteq\{1,2,\dots,n\}} \hat{\phi}(S)\chi_S\right)\left(\sum_{T\subseteq\{1,2,\dots,n\}} \hat{\psi}(T)\chi_T\right) \\ &= \sum_{S,T\subseteq\{1,2,\dots,n\}} \hat{\phi}(S)\hat{\psi}(T)\chi_{(S\setminus T)\cup(T\setminus S)}.\end{aligned}$$

Applying Proposition 2.2,

$$\|\!|\phi \cdot \psi\|\!| \leq \sum_{S,T\subseteq\{1,2,\dots,n\}} |\hat{\phi}(S)| \, |\hat{\psi}(T)|.$$

The right-hand side of this inequality is clearly $\|\!|\phi\|\!| \, \|\!|\psi\|\!|$. □

We will use the norm $\|\!|\cdot|\!\|$ in conjunction with the operator L_k to refer to the sum of the absolute values of the Fourier coefficients of given order k:

$$|\!|\!| L_k\,\phi |\!|\!| = \sum_{S\in\mathcal{P}_{n,k}} |\hat{\phi}(S)|.$$

2.6. Sign-representations

Let $f\colon X \to \{0,1\}$ be a given Boolean function, for a finite subset $X \subset \mathbb{R}^n$. The *threshold degree* of f, denoted $\deg_{\pm}(f)$, is the least degree of a real polynomial p that represents f in sign: $\operatorname{sgn} p(x) = (-1)^{f(x)}$ for each $x \in X$. The term "threshold degree" appears to be due to Saks [**108**]. Equivalent terms in the literature include "strong degree" [**10**], "voting polynomial degree" [**81**], "polynomial threshold function degree" [**97**], and "sign degree" [**32**]. The AND_n function has threshold degree 1, as the polynomial

$$p(x) = -\sum_{i=1}^{n} x_i - \frac{1}{2} + n$$

represents it in sign.

A closely related notion of the threshold degree is the *approximate degree*, for which we are interested in polynomials that approximates f. In particular, for any $\varepsilon \in [0, 1/2)$ a polynomial p is said to ε-approximate f if

$$|f(x) - p(x)| \leq \varepsilon, \qquad\qquad \forall\, x \in X.$$

The ε-approximate degree of f, denoted $\deg_{\varepsilon}(f)$, is the minimum degree of a polynomial p that ε-approximates f. Clearly, $\deg_{\varepsilon}(f)$ is nonincreasing in ε for $\varepsilon \in [0, 1/2)$, and

$$\deg_{\pm}(f) = \lim_{\varepsilon \nearrow \frac{1}{2}} \deg_{\varepsilon}(f).$$

It is also not hard to show that for any constant $0 < \delta < \varepsilon < 1/2$, $\deg_\varepsilon(f) = \Theta(\deg_\delta(f))$, as essentially what we need is an S-shape polynomial of constant degree that maps $[0, \varepsilon] \mapsto [0, \delta]$, and $[1-\varepsilon, 1] \mapsto [1-\delta, 1]$. The following simple and ingenious candidate is due to Buhrman et al. [**31**],

$$B_d(t) = \sum_{i=\lceil \frac{d+1}{2} \rceil}^{d} \binom{d}{i} t^i (1-t)^{d-i}.$$

This family of polynomials calculates the probability that there are more heads than tails when flipping d random coins where each coin has independent probability of t to be head. By Chernoff bound, $B_d : [0, \varepsilon] \mapsto [0, 2^{-\Theta_\varepsilon(d)}]$. In particular, take $d = O(-\ln\delta/(1-2\varepsilon)^2)$, then the polynomial $B_d(p)$ δ-approximates f where p is any polynomial that ε-approximates f.

One of the first results on polynomial representations of Boolean functions was the following tight lower bound on the threshold degree of the Minsky-Papert function [**90**]. The generalized *Minsky–Papert function* $\mathrm{MP}_{m,r} \colon (\{0,1\}^r)^m \to \{0,1\}$ is given by $\mathrm{MP}_{m,r}(x) = \bigwedge_{i=1}^{m} \bigvee_{j=1}^{r} x_{i,j}$. We abbreviate $\mathrm{MP}_m = \mathrm{MP}_{m,m^2}$, which is the right setting of parameters for most of our applications.

Theorem 2.4 (Minsky and Papert). $\deg_\pm(\mathrm{MP}_m) = \Omega(m)$.

Three new proofs of this lower bound, unrelated to Minsky and Papert's original proof, were discovered recently in [**122**]. Threshold degree admits the following dual characterization, obtained by appeal to linear programming duality.

Fact 2.5. *Let $f\colon X \to \{0,1\}$ be a given Boolean function on a finite subset X of Euclidean space. Then $\deg_{\pm}(f) \geq d$ if and only if there exists $\psi\colon X \to \mathbb{R}$ such that*

$$\begin{aligned} &(-1)^{f(x)}\psi(x) \geq 0, && x \in X, \\ &\langle \psi, P\rangle = 0, && \forall\, P : \deg P < d, \\ &\psi \not\equiv 0. \end{aligned}$$

The function ψ acts as a *witness* for the threshold degree of f, and is called a *dual polynomial* due to its origin in a dual linear program. Here we present a simple proof for the case when X is $\{0,1\}^n$, using Gordan's theorem [**67**].

Theorem 2.6 (Gordan's Theorem). *For any matrix $M \in \mathbb{R}^{m\times n}$, exactly one of the following is true*

(i) $Mx = 0$, *for some nonzero $x \in \mathbb{R}^n$ and $x \geq 0$,*

(ii) $y^T A > 0$, *for some $y \in \mathbb{R}^m$.*

To apply Gordan's Theorem, we translate $\deg_{\pm}(f) \geq d$ into the linear algebra language. Note $\{\chi_S : |S| < d\}$ is a basis for the vector space spanned by polynomials of degree less than d. Consider matrix $M \in \mathbb{R}^{\binom{[n]}{<d}\times\{0,1\}^n}$, where

$$M = [(-1)^{f(x)} \cdot \chi_S(x)]_{S\in\binom{[n]}{<d}, x\in\{0,1\}^n}.$$

No sign representing polynomial P for f of degree less than d, is equivalent to say that no $y = (y_S)_{S\in\binom{[n]}{<d}}$ satisfies

$$y^T M > 0.$$

Therefore by Gordan's Theorem, $\deg_\pm(f) \geq d$ if and only if there is $\phi : \{0,1\}^n \to \mathbb{R}$, $\phi \gneq 0$, and

$$\sum_{x\in\{0,1\}^n} \phi(x)(-1)^{f(x)}\chi_S(x) = 0, \qquad\qquad \forall\, S : |S| < d.$$

Take $\psi = (-1)^f\phi$, then Fact 2.5 is true. We refer the reader to [**10, 97, 119**] for a detailed proof of Fact 2.5. The following equivalent statement is occasionally more convenient to work with.

FACT 2.7. *For every Boolean function $f\colon X \to \{0,1\}$ on a finite subset X of Euclidean space,*

$$\deg_\pm(f) = \max_{\mu:\text{distribution on } X} \ \min_{\substack{\text{polynomial } P:\\ \langle(-1)^f\cdot\mu),P\rangle\neq 0}} \deg P. \tag{2.6.1}$$

We now define a generalization of threshold degree inspired by the dual view in Fact 2.7. For a function $f\colon X \to \{0,1\}$ and a real number $0 \leq \gamma \leq 1$, let

$$\deg_\pm(f,\gamma) = \max_{\substack{\mu:\text{distribution on } X\\ \mu\geq\gamma/|X| \text{ on } X}} \ \min_{\substack{\text{polynomial } P:\\ \langle(-1)^f\cdot\mu),P\rangle\neq 0}} \deg P. \tag{2.6.2}$$

We call this quantity the *γ-smooth threshold degree of f*, in reference to the fact that the maximization in (2.6.2) is over probability distributions μ that place on every point of the domain at least a γ fraction of the weight the point would receive under the uniform distribution. A glance at (2.6.1) and (2.6.2) reveals that $\deg_\pm(f,\gamma)$ is monotonically nonincreasing in γ, with the limiting case $\deg_\pm(f,0) = \deg_\pm(f)$.

FACT 2.8. *For every nonconstant function $f\colon X \to \{0,1\}$,*

$$\deg_\pm\left(f, \frac{1}{2}\right) \geq 1.$$

Proof. Define $\mu = \frac{1}{2}\mu_0 + \frac{1}{2}\mu_1$, where μ_i be the uniform probability distribution on $f^{-1}(i)$. Then clearly $\langle(-1)^f \cdot \mu), c\rangle = 0$ for any nonzero constant c and $\mu \geq \frac{1}{2}\max\{\mu_0, \mu_1\} \geq \frac{1}{2|X|}$ on X. □

2.7. Query complexity

For a Boolean function $f : \{0,1\}^n \to \{0,1\}$, the query complexity of f measures the number of bits that we need to know algorithmically to determine the value of f. Next, we give concrete definitions of the query complexity in different models.

2.7.1. Deterministic query complexity.

A deterministic query algorithm of function f is a *decision tree.* A decision tree is a binary tree T in which each leaf is labeled with 0 or 1, and each inner node is labeled with some variable x_i. Given an input $x \in \{0,1\}^n$, the algorithm first query the variable labeled in the root. Depending on the outcome of the query, the algorithm moves from the root to either the left child or the right child, and query the variable in the corresponding child. The algorithm repeats this procedure till reaching a leaf and outputs the label on that leaf. The *depth* of a decision tree T, is the number of edges from the root to the furthest leaf. We say T computes f if T outputs the correct value $f(x)$ for all $x \in \{0,1\}^n$. The deterministic query complexity of f, denoted $D^{\mathrm{dt}}(f)$, is the minimum depth of a decision tree T that computes f. The superscript dt stands for "decision tree," and are used to distinguish from the deterministic communication complexity. Clearly $0 \leq D^{\mathrm{dt}}(f) \leq n$. For example $D^{\mathrm{dt}}(\mathrm{AND}_n) = n$, since when the input $x = 1^n$, any decision tree has to query all the variables to be sure there is no 0.

The above definition is for functions with domain $\{0,1\}^n$ and range $\{0,1\}$. But it is clear that a similar definition works for function $f : X_1 \times X_2 \times \cdots \times X_n \to Y$, for arbitrary finite sets $X_1, X_2, \ldots, X_n$ and an arbitrary set Y.

2.7.2. Randomized query complexity. Now we allow decision trees to access random bits. There are different ways to define randomized decision trees. We adopt the randomized decision trees with bounded-error, which is the counterpart of Monte Carlo algorithms for Turing machines. In particular, a randomized decision tree is a distribution μ on decision trees. Its depth is the largest depth of any decision tree $T \in \mathrm{supp}(\mu)$. An ε-error randomized decision tree for f outputs a correct answer with probability at least $1 - \varepsilon$ on any input, i.e.,

$$\Pr_{T \sim \mu}[T(x) = f(x)] \geq 1 - \varepsilon, \qquad \forall\, x \in \{0,1\}^n.$$

The randomized query complexity of f, denoted $R^{\mathrm{dt}}_{\varepsilon}(f)$, is the minimum depth of an ε-error randomized decision tree for f. We omit ε when $\varepsilon = 1/3$.

Clearly, for any $\varepsilon \geq 0$ and function f, $R^{\mathrm{dt}}_{\varepsilon}(f) \leq D^{\mathrm{dt}}(f)$. To see that the randomized decision trees can be strictly more powerful, consider the MAJ_3 problem. $D^{\mathrm{dt}}(\mathrm{MAJ}_3) = 3$, but $R^{\mathrm{dt}}(\mathrm{MAJ}_3) = 2$. For the randomized decision tree, consider the following algorithm:

(i) Take two random bits x_i, x_j (without repetition);

(ii) If $x_i = x_j$, output x_i; otherwise output a random value of 0 or 1.

This algorithm makes no error for inputs 000 or 111. For input 110 (or any permutations of 110), the algorithm outputs the correct answer with probability 2/3. To exhibit a larger gap, a standard trick is to take function composition. Let

$$\mathrm{MAJ}^{(i)} = \mathrm{MAJ}^{(i-1)} \circ \mathrm{MAJ}_3, \qquad i = 2, 3, \ldots$$
$$\mathrm{MAJ}^{(1)} = \mathrm{MAJ}_3.$$

A similar argument can be used to show that

$$R^{\mathrm{dt}}(\mathrm{MAJ}^{(d)}) = O((D^{\mathrm{dt}}(\mathrm{MAJ}^{(d)}))^{0.8927\ldots}.$$

For total functions $f : \{0,1\}^n \to \{0,1\}$, it is well-known that randomized and deterministic query complexity are polynomially related due to Nisan [**92**].

$$D^{\mathrm{dt}}(f) \le 27(R^{\mathrm{dt}}(f))^3.$$

Open Problem 2.9. What is the largest separation between D^{dt} and R^{dt}?

For partial functions $f : \{0,1\}^n \to \{0,1,*\}$ though, things are quite different. This is a common theme in the query world. Consider the gap majority problem

$$\mathrm{GapMAJ}_{3n}(x) = \begin{cases} 0 & \|x\|_1 \le n, \\ 1 & \|x\|_1 \ge 2n, \\ * & \text{otherwise.} \end{cases}$$

Then $D^{\mathrm{dt}}(\mathrm{GapMAJ}_{3n}) = \Omega(n)$, as only when we see at least $(n+1)$ 0s or $(n+1)$ 1s, we can safely determine the value of the gap majority problem. But

$$R^{\mathrm{dt}}(\mathrm{GapMAJ}_{3n}) = 1,$$

i.e., query a random bit and output its value.

2.7.3. Quantum query complexity.

Finally, we review the quantum query algorithm. Before that, we first briefly discuss some basics of a quantum system. The simplest quantum system consists of a single quantum bit. Its state is described by a unit vector in $\mathbb{C}^2$ with orthonormal basis $|0\rangle$ and $|1\rangle$, corresponding to the classical state of bit 0 and bit 1. In particular, the quantum bit in *superposition* of bit 0 and bit 1 is represented by a unit column vector

$$|\phi\rangle = a|0\rangle + b|1\rangle,$$

for $a, b \in \mathbb{C}$, such that $|a|^2 + |b|^2 = 1$. The *amplitudes* a and b have the classical interpretation that when make a measurement of the quantum bit, with probability of $|a|^2$ the bit 0 will be observed and with probability of $|b|^2$ the bit 1 will be observed. More generally, an m-bit quantum system is described by $\mathbb{C}^{2^m}$ with basis

$$\{|x\rangle : x \in \{0,1\}^m\}.$$

The elements in the basis correspond to the set of classical state of m bits. The m quantum bits in superposition are described by a unit vector

$$\sum_{x\in\{0,1\}^m} a_x|x\rangle,$$

and $\|a\|_2 = 1$

A quantum query algorithm uses a quantum system described by tensor product $\mathcal{I} \otimes \mathcal{W}$, where the $\mathcal{I} = \mathbb{C}^{\lceil \log n \rceil}$ corresponds to the query indices, and $\mathcal{W}$ corresponds to the workspace. So its state can be described by

$$|\psi\rangle = \sum_{i=1}^{n} a_i|i, w_i\rangle,$$

where $|w_i\rangle$ is a unit vector in $\mathcal{W}$. To make a query, apply the query operator O, which is a unitary transformation, defined as follows[2],

$$O : |i, w_i\rangle \mapsto (-1)^{x_i}|i, w_i\rangle.$$

A k query algorithm is the following sequence of operators applied on the initial state $|0, 0\rangle$ (note that 0 here is the 0 vector in the corresponding spaces),

$$U_0, O, U_1, O, \ldots, O, U_k,$$

[2] The effect of the operator O on the entire vector space can be generalized by linearity.

where U_i can be arbitrary unitary transformations, corresponding to quantum computations. After the running the algorithm, the final state is

$$|\psi\rangle = U_k O \cdots O U_1 O U_0 |0, 0\rangle.$$

To output, measure the last bit of $|\psi\rangle$.

A quantum query algorithm is said to compute f with error ε, if it outputs the correct answer with probability at least $1 - \varepsilon$ on any input. The ε-error quantum query complexity of f, denoted $Q^{\mathrm{dt}}_{\varepsilon}(f)$, is the minimum number of queries of a quantum query algorithm that computes f with error ε. We omit ε for $\varepsilon = 1/3$.

Quantum query algorithm is more powerful than the classical query algorithm. The well-known Grover's search algorithm shows that $Q^{\mathrm{dt}}(\mathrm{AND}_n) = O(\sqrt{n})$, but $R^{\mathrm{dt}}(\mathrm{AND}_n) = \Omega(n)$. Like many other measures in the query model, Q^{dt} and R^{dt} are also polynomially related for total functions. For partial functions, they can be arbitrarily separated. We will discuss the relationship between quantum and classical query complexity in more detail in Chapter 6.

CHAPTER 3

Communication protocols and communication complexity

In this chapter, we define the communication protocol and provide a short introduction to the communication complexity. We will only cover the key definitions and notations. A classical and excellent reference on communication complexity is the monograph by Kushilevitz and Nisan [**84**]. A coverage on the more recent development of this field is beautifully presented by Rao and Yehudayoff [**102**].

3.1. Communication models

3.1.1. Deterministic communication. We adopt the standard two-party model of *deterministic* communication due to Yao [**139**]. In this model, Alice and Bob receive inputs $x \in X$ and $y \in Y$, respectively, where X and Y are some finite sets fixed in advance. Their goal is to compute a given function $f : X \times Y \to \{0, 1\}$. They communicate by sending each other symbols from a fixed alphabet Σ. The most common alphabet is $\Sigma = \{0, 1\}$, but we will encounter others as well. The communication between Alice and Bob is governed by an agreed-upon *protocol* π. At any given time, the protocol specifies, based on the sequence of symbols exchanged so far between Alice and Bob, whether the communication is to continue and if so, who should send the next symbol. This next symbol is also specified by the protocol, based on the sender's input as well as the sequence of symbols exchanged so far between Alice and Bob. Denote the entire communication history of the protocol running by $\pi(x, y)$. The *communication cost* of the protocol π, denoted $|\pi|$, is the worst-case number of transmissions, i.e., $|\pi| = \max_{x,y} |\pi(x, y)|$. The *output* is determined by the entire communication history $\pi(x, y)$. In particular, the protocol π has a interpretation function $\mathrm{out} : \Sigma^* \to \{0, 1\}$, and the output of π is $\mathrm{out}(\pi(x, y))$. We say that a protocol π computes f, if $\mathrm{out}(\pi(x, y)) = f(x, y)$ for all $(x, y) \in X \times Y$. The communication complexity of f, is the minimum communication cost of a protocol that computes f. Consider the classical example, the *equality* problem,

$$\mathrm{EQ}(x, y) = \begin{cases} 1 & x = y, \\ 0 & x \neq y. \end{cases}$$

A trivial protocol is the following: Alice sends Bob x, and Bob announces 1 if and only if $x = y$. This is in fact the best we can do deterministically, thus,

$$D(\mathrm{EQ}) = n + 1. \tag{3.1.1}$$

A protocol π is said to be *in canonical form* if the following two conditions hold: (i) the number of symbols exchanged between Alice and Bob is an even integer and is the same for all inputs $x \in X$ and $y \in Y$; (ii) Alice and Bob take turns sending each other one symbol at a time, with Alice sending the first symbol. A moment's thought reveals that any protocol π can be simulated by a protocol in canonical form with the same alphabet and at most double the communication cost.

A communication protocol π over alphabet Σ can be visualized in terms of a regular tree of depth $|\pi|$, called the *protocol tree.* Every internal vertex of the protocol tree has precisely $|\Sigma|$ outgoing edges, each labeled with a distinct symbol of the alphabet. A vertex of the protocol tree corresponds in a one-to-one manner to a state of the protocol at some point in time. Specifically, the vertex reachable from the root via the path $v \in \Sigma^*$ corresponds to the point in time when the symbols exchanged between Alice and Bob so far are precisely $v_1, v_2, \ldots, v_{|v|}$, in that order. In particular, the root vertex corresponds to the point in time just before the communication starts, and a leaf corresponds to a point in time when the communication has ended. Every internal vertex of the protocol tree is said to be *owned* by either Alice or Bob, corresponding to the identity of the speaker at that point in time. For a given input $x \in X$, the protocol specifies a unique outgoing edge for every vertex owned by Alice, corresponding to the symbol that she would send at that point in time with x as her input. Analogously, for any $y \in Y$, the protocol specifies a unique outgoing edge for every vertex owned by Bob. On any input pair x, y, Alice and Bob's edges determine a unique root-to-leaf path. Execution of the protocol corresponds to a walk down this unique root-to-leaf path defined by Alice and Bob's edges, and the output of the protocol is the label on the corresponding leaf. Adopting this view of communication, we will henceforth identify Alice's input with a set of edges, one for each vertex that Alice owns; and likewise for Bob. Observe that if the protocol is in canonical form, Alice and Bob's inputs are a set of outgoing edges for the even-depth vertices and a set of outgoing edges for the odd-depth vertices, respectively, one such edge per vertex.

Finally, we make a remark on the key combinatorial property of a deterministic protocol first observed by Yao [**139**].

PROPOSITION 3.1. *For any $x, x' \in X$, $y, y' \in Y$, if $\pi(x,y) = \pi(x',y')$. Then,*

$$\pi(x',y) = \pi(x,y') = \pi(x,y).$$

We refer to $A \times B$ a *combinatorial rectangle* or simply a *rectangle* for $A \subseteq X, B \subseteq Y$. The above proposition implies that a deterministic protocol π partitions $X \times Y$ into at most $2^{|\pi|}$ rectangles according to all possible communication transcripts.

Sometimes the inputs are sampled from a prior distribution μ. In this case, it is reasonable to allow the protocol to err in a small fraction of the inputs. We can define distributional complexity of f with respect to distribution μ as follows

$$D^{\mu}_{\varepsilon}(f) = \min\{|\pi| \mid \exists\, \mathrm{out} : \Sigma^* \to \{0,1\}, \\ \Pr_{(x,y)\sim\mu}[\mathrm{out}(\pi(x,y)) = f(x,y)] \geq 1 - \varepsilon\}.$$

3.1.2. Randomized communication. In the *randomized* model, the parties are provided with their own private source of unlimited random bits. The communication between Alice and Bob is governed by an agreed-upon protocol π and the outcome of the random strings r_a and r_b owned by Alice and Bob, respectively. In particular, at any given time, the protocol π specifies, based on the sequence of symbols exchanged so far between Alice and Bob, whether the communication is to continue and if so, who should send the next symbol. This next symbol is specified by the protocol, based on the sender's input, the sender's random string, and the sequence of symbols exchanged so far. An *ε-error protocol* for f is one which on every input pair (x, y), produces the correct answer $f(x, y)$ with probability at least $1 - \varepsilon$. The *ε-error randomized communication complexity* of f, denoted $R_\varepsilon(f)$, is the least cost of an ε-error randomized protocol for f. Since we can repeat a protocol for multiple times using fresh random strings, and take the majority vote, it is immediate that for any two constant $0 < \varepsilon < \delta < 1/2$, $R_\varepsilon(f) = \Theta(R_\delta(f))$. This simple fact releases us from worrying about the exact error parameter as long as it is a constant in $(0, 1/2)$. We use $R(f)$ to denote $R_{1/3}(f)$, where $1/3$ is chosen for aesthetic reason.

A variance of the randomized model is the public randomized communication model. The main difference is that in the public randomized communication model, Alice and Bob have a *shared* source of unlimited random bits. In this model, the shared random string r replaces the role of the private random strings in the former model. In particular, the next symbol is specified based on the shared random string r, sender's input and the communication history so far. A moment thought reveals that a public random protocol can be viewed as a distribution on deterministic protocols. It is clear that a protocol with public random bits can easily simulate a protocol with private random sources without any effort. But it is not so obvious whether it is true the other way around. Somewhat surprisingly, shared randomness has essentially no effect on the randomized communication complexity $R(f)$ for any f, due to the well-known result by Newman [**91**]. Let $R_\varepsilon^{\text{pub}}(f)$ be the communication complexity of the public randomized communication complexity, Newman proved the following relationship between $R_\varepsilon^{\text{pub}}(f)$ and $R_\varepsilon(f)$.

THEOREM 3.2 (Newman). *Let $f : \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ be any Boolean function. For every $\varepsilon, \delta > 0$,*

$$R_{\varepsilon+\delta}(f) \leq R_\varepsilon^{\text{pub}}(f) + O\left(\log n + \log \frac{1}{\delta}\right).$$

The above theorem is very general. In essence, it proves that for any probabilistic computation, $O(\log n + \log 1/\delta)$ number of random bits is sufficient at the cost of increasing the error probability by at most δ (at a cost of computation).

An randomized protocol can be much more efficient than that of a deterministic protocol. We take the equality function as an example. Consider the following protocol:

(i) Alice and Bob use public randomness to sample $z, z' \in \{0,1\}^n$,

(ii) Interpreting x, y, z, z' as vectors in $\mathbb{F}_2^n$, Alice sends $\langle x, z\rangle$ and $\langle x, z'\rangle$ to Bob,

(iii) Bob outputs 1 if and only if $\langle x, z\rangle = \langle y, z\rangle$ and $\langle x, z'\rangle = \langle y, z'\rangle$.

It is easy to verify that if $x = y$, then Bob always outputs 1. Otherwise Bob outputs 1 with probability of $1/4$.

Yao made the following astonishing observation that connects the randomized communication complexity and distributional communication complexity.

Theorem 3.3 (Yao's principle). $R_\varepsilon^{\text{pub}}(f) = \max_\mu D_\varepsilon^\mu(f)$.

Finally, we mention the following proposition that generalizes Proposition 3.1 for the randomized protocols.

Proposition 3.4. *Let $\mathcal{T}$ be the set of all possible transcripts of protocol π. There are functions $p_x : \mathcal{T} \to [0,1]$, $q_y : \mathcal{T} \to [0,1]$ for all $x \in X$ and $y \in Y$, such that*

$$\Pr[\pi(x,y) = \tau] = p_x(\tau) q_y(\tau), \qquad \forall\, \tau \in \mathcal{T}.$$

3.1.3. Unbounded-error communication. A particularly interesting situation is when the randomized protocols have error probability close to that of random guessing, i.e., 1/2. There are two natural ways to define the communication complexity of a problem f in this setting. The *communication complexity of f with unbounded error*, introduced by Paturi and Simon [**101**], is the quantity

$$\mathsf{UPP}(F) = \min_{0<\varepsilon<1/2} R_\varepsilon(F). \tag{3.1.2}$$

Here, the error is unbounded in the sense that it can be arbitrarily close to 1/2. To emphasize this difference of allowed error, we sometimes refer to the the communication complexity with constant error as the bounded-error communication complexity. In contrast to the bounded-error communication model discussed in the previous section, it is important here that Alice and Bob have access only to their private random sources. Otherwise they can solve any problem exchanging 2 bits by the following protocol: Alice sends 1 if the public random string $r = x$, sends 0 otherwise. Bob outputs a uniformly random bit if $r \neq x$, otherwise he sends $f(r, y)$. Clearly, this protocol has success probability better than 1/2, and costs only 2 bits. The unbounded-error communication occupies a special place in the study of communication because it is more powerful than almost any other standard model, e.g., the randomized communication model and the quantum communication model. Unbounded-error protocols represent a frontier in communication complexity theory in that they are the most powerful protocols for which explicit lower bounds are currently known.

Babai et al. [**11**] proposed an alternate quantity, which includes an additive penalty term that depends on the error probability:

$$\mathsf{PP}(f) = \min_{0<\varepsilon<1/2} \left\{ R_\varepsilon(f) + \log \frac{1}{\frac{1}{2}-\varepsilon} \right\}. \tag{3.1.3}$$

This quantity is known as the *communication complexity of f with weakly unbounded error.* It is clear that

$$1 \leq \mathsf{UPP}(f) \leq \mathsf{PP}(f) \leq n+1$$

for every communication problem f, with an exponential gap achievable between the two complexity measures [**32, 113**]. It is also clear that

$$\mathsf{PP}(f) \leq R(f) + 3.$$

An exponential gap is achieved by the *disjointness* problem

$$\mathrm{DISJ}_n(x,y) = \neg \bigvee_{i=1}^{n} (x_i \wedge y_i).$$

The celebrated result due to Kalyanasundaram and Schnitger [**75**] proves that $R(\mathrm{DISJ}_n) = \Omega(n)$. On the other hand, it is not hard to see that $\mathsf{PP}(\mathrm{DISJ}_n) = O(\log n)$.

3.1.4. Quantum communication. The *quantum communication complexity* was first introduced by Yao [**140**]. As before, one considers the problem of solving a function $f : X \times Y \to \{0,1\}$, with the inputs distributed between Alice and

Bob. But now, Alice and Bob can exchange quantum bits, and take advantage of quantum *entanglement*. There are various ways to define quantum communication protocols. In this dissertation, we follow closely Razborov's description [**105**]. We will be working with the tensor product $\mathcal{A} \otimes \mathcal{C} \otimes \mathcal{B}$, where $\mathcal{A}, \mathcal{B}, \mathcal{C}$ are complex vector spaces. $\mathcal{A}, \mathcal{B}, \mathcal{C}$ represent Alice's workspace, Bob's workspace, and the communication channel respectively. $\mathcal{C}$ is a dimension 2 vector space with basis $\{|0\rangle, |1\rangle\}$. Initially, the quantum system without prior entanglement has state

$$|\phi_0\rangle = |x, 0\rangle|0\rangle|y, 0\rangle,$$

while the quantum system with prior entanglement has state

$$|\phi_0\rangle = \frac{1}{\sqrt{|E|}} \sum_{e \in E} |x, 0, e\rangle|0\rangle|y, 0, e\rangle,$$

Here, $|x, 0, e\rangle \in \mathcal{A} = \mathcal{X} \otimes \mathcal{W} \otimes \mathcal{E}$ and $|y, 0, e\rangle \in \mathcal{B} = \mathcal{Y} \otimes \mathcal{W} \otimes \mathcal{E}$, where $\mathcal{X}$ and $\mathcal{Y}$ are the complex vector spaces describing the input sets X and Y, $\mathcal{W}$ corresponds to the auxiliary quantum bits and $\mathcal{E}$ corresponds to the prior entanglements. A quantum communication protocol can be described by a sequence of unitary operators acting on $\mathcal{A} \otimes \mathcal{C} \otimes \mathcal{B}$:

$$U_1 \otimes I_{\mathcal{B}}, I_{\mathcal{A}} \otimes U_2, U_3 \otimes I_{\mathcal{B}}, \ldots, I_{\mathcal{A}} \otimes U_{2k},$$

where $I_{\mathcal{A}}, I_{\mathcal{B}}$ are identity transformation in $\mathcal{A}$ and $\mathcal{B}$, respectively, and U_i is a unitary transformation in $\mathcal{A} \otimes \mathcal{C}$ for odd i and a unitary transformation in $\mathcal{C} \otimes \mathcal{B}$ for even i. The transformation U_i corresponds to Alice (or Bob) making local computation and then sending a quantum bit to Bob (or Alice). In the end of the protocol, the quantum state is

$$|\phi_{2k}\rangle = (I_{\mathcal{A}} \otimes U_{2k}) \cdots (I_{\mathcal{A}} \otimes U_2)(U_1 \otimes I_{\mathcal{B}})|\phi_0\rangle.$$

The cost of the protocol is the length of the sequence, e.g., $2k$ here. To output, measure $\mathcal{C}$. Mathematically, let v be the projection of $|\phi_{2k}\rangle$ onto the subspace $\mathcal{A} \otimes |1\rangle \otimes \mathcal{B}$. Then the protocol outputs 1 with probability $\|v\|^2$, and outputs 0 with probability $1 - \|v\|^2$. The ε-error quantum communication complexity without prior entanglement of function f, denoted $Q_\varepsilon(f)$, is the cost of a quantum protocol that outputs a correct answer with probability at least $1 - \varepsilon$ on any input x, y. We abbreviate $Q(f) = Q_{1/3}(f)$. The corresponding quantum communication complexity with prior entanglement is denoted by $Q^*_\varepsilon(f)$.

It is clear that $Q_\varepsilon(f) \leq R_\varepsilon(f)$ for any function f and ε, as the quantum bits can be used as private randomness. Quantum protocols can be strictly more powerful. For example, the quantum communication complexity of the disjointness problem is $\Theta(\sqrt{n})$ [**105, 1**], while we have discussed that the disjointness problem has linear randomized communication complexity. However, how much more power a quantum protocol can have is not completely understood. It is a major open problem to determine if the quantum communication complexity and the randomized communication complexity are polynomially related for total functions. A detailed discussion on this issue is deferred to Chapter 6.

3.1.5. Multiparty communication. The two-party communication models have analogues for three or more parties. In this dissertation, we will only be interested in the *number-on-forehead* model of multiparty communication, introduced by Chandra et al. [**40**]. The model features ℓ communicating players, tasked with computing a Boolean function $F\colon X_1 \times X_2 \times \cdots \times X_\ell \to \{0,1\}$ for some finite sets $X_1, X_2, \ldots, X_\ell$. A given input $(x_1, x_2, \ldots, x_\ell) \in X_1 \times X_2 \times \cdots \times X_\ell$ is distributed among the players by placing x_i, figuratively speaking, on the forehead of the ith player (for $i = 1, 2, \ldots, \ell$). In other words, the ith player knows the arguments

$x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_\ell$ but not x_i. The players communicate by sending broadcast messages, taking turns according to a protocol agreed upon in advance.

In a randomized protocol, each of the players privately holds an unlimited supply of uniformly random bits, which he can use along with his available arguments when deciding what message to send at any given point in the protocol. The players' objective is to compute $F(x_1, x_2, \ldots, x_\ell)$. An *$\varepsilon$-error protocol* for F is one which, on every input $(x_1, x_2, \ldots, x_\ell)$, produces the correct answer $F(x_1, x_2, \ldots, x_\ell)$ with probability at least $1 - \varepsilon$. The *cost* of a protocol is the total bit length of the messages broadcast by all the players in the worst case.[1] The *ε-error randomized communication complexity* of F, denoted $R_\varepsilon(F)$, is the least cost of an ε-error randomized protocol for F. As a special case of this model for $\ell = 2$, one recovers the original two-party model.

As combinatorial rectangles are fundamental objects for two-party protocols, the counterpart in the multiparty protocols is *cylinder intersections.* An *ℓ-dimensional cylinder intersection* is a function $\chi\colon X_1 \times X_2 \times \cdots \times X_\ell \to \{0, 1\}$ of the form

$$\chi(x_1, x_2, \ldots, x_\ell) = \prod_{i=1}^{\ell} \chi_i(x_1, \ldots, x_{i-1}, x_{i+1}, \ldots, x_\ell),$$

where $\chi_i\colon X_1 \times \cdots \times X_{i-1} \times X_{i+1} \times \cdots \times X_\ell \to \{0, 1\}$. In other words, an ℓ-dimensional cylinder intersection is the product of ℓ functions with range $\{0, 1\}$, where the ith function does not depend on the ith coordinate but may depend arbitrarily on the other $\ell - 1$ coordinates. Introduced by Babai et al. [**12**], cylinder intersections are the fundamental building blocks of communication protocols and for that reason play a central role in the theory. The 2-dimensional cylinder is a rectangle. A deterministic protocol partitions the input space $X_1 \times X_2 \times \cdots \times X_\ell$ into cylinder intersections.

[1] The contribution of a b-bit broadcast to the protocol cost is b rather than $\ell \cdot b$.

3.2. Communication complexity lower bounds

In this section, we discuss several techniques to show strong communication lower bounds. A common theme under these techniques is to relate the communication complexity to some other measures easier to analyze. There are many other influential techniques to prove communication lower bounds like the rectangle bounds and information complexity that are out of the scope of this dissertation. Interested readers are referred to [**86, 102**].

3.2.1. Query-to-communication lifting. Query complexity are usually much easier to analyze than the communication complexity. A powerful tool developed in recent years transforms query complexity lower bounds to communication complexity lower bounds in various models [**64, 44, 65, 43**]. The intuition behind this tool is very simple. Let $f : \{0,1\}^n \to \{0,1\}$ be any Boolean function with large query complexity. Consider a small-size gadget $g : \{0,1\}^\ell \times \{0,1\}^\ell \to \{0,1\}$, and the communication problem $f \circ g : \{0,1\}^{n\ell} \times \{0,1\}^{n\ell} \to \{0,1\}$,

$$f \circ g(x, y) = f(g(x_1, y_1), g(x_2, y_2), \ldots, g(x_n, y_n)).$$

A naïve communication protocol is to simulate the query algorithm for f: whenever the algorithm queries the ith variable, Alice and Bob compute $g(x_i, y_i)$ by communicating at most ℓ bits. Thus,

$$D(f \circ g) = O(\ell \cdot D^{\mathrm{dt}}(f)),$$

$$R(f \circ g) = O(\ell \cdot R^{\mathrm{dt}}(f)).$$

For a gadget g complicated enough that hides $g(x_i, y_i)$ well, we expect that there is no better protocol than the naïve protocol, thus,

$$D(f \circ g) = \Omega(\ell \cdot D^{\mathrm{dt}}(f)),$$
$$R(f \circ g) = \Omega(\ell \cdot R^{\mathrm{dt}}(f)).$$

Of course, this cannot be true for every gadget g. For example, if f and g are both the parity function, then the communication complexity of $\bigoplus_{i=1}^{n}(x_i \oplus y_i)$ is only a constant, but the query complexity of the parity function is n.

Ideally, we want g to be a constant-size gadget. However, whether there is such a gadget is still open. A well-studied function g in this line of research is the inner product function $\mathrm{IP}_\ell\colon \{0,1\}^\ell \times \{0,1\}^\ell \to \{0,1\}$, given by $\mathrm{IP}_\ell(u,v) = \bigoplus_{i=1}^{\ell}(u_i \wedge v_i)$. In particular, Chattopadhyay, Filmus, Koroth, Meir, and Pitassi [**43**, Theorem 1] prove that

$$R_{1/3}(f \circ \mathrm{IP}_{c\log n}) = \Omega(R^{\mathrm{dt}}_{1/3}(f)\log n) \tag{3.2.1}$$

for every (possibly partial) function f on $\{0,1\}^n$, where $c > 1$ is an absolute constant.

Open Problem 3.5. Find a constant-size gadget $g : \{0,1\}^c \times \{0,1\}^c$, such that for any $f : \{0,1\}^n \to \{0,1\}$,

$$D(f \circ g) = \Omega(D^{\mathrm{dt}}(f)).$$

3.2.2. Discrepancy and Sign-rank.

For a Boolean function $F\colon X_1 \times X_2 \times \cdots \times X_\ell \to \{0,1\}$ and a probability distribution μ on $X_1 \times X_2 \times \cdots \times X_\ell$, the *discrepancy of F with respect to μ* is given by

$$\mathrm{disc}_\mu(F) = \max_{\chi} \left| \sum_{x \in X_1 \times X_2 \times \cdots \times X_\ell} (-1)^{F(x)} \mu(x) \chi(x) \right|,$$

where the maximum is over cylinder intersections χ. The minimum discrepancy over all distributions is denoted

$$\operatorname{disc}(F) = \min_{\mu} \operatorname{disc}_{\mu}(F).$$

The *discrepancy method* [**47, 12, 84**] is a classic technique that bounds randomized communication complexity from below in terms of discrepancy.

THEOREM 3.6 (Discrepancy method). *Let $F\colon X_1 \times X_2 \times \cdots \times X_\ell \to \{0,1\}$ be a given communication problem. Then*

$$2^{R_\varepsilon(F)} \geq \frac{1-2\varepsilon}{\operatorname{disc}(F)}.$$

This theorem is an immediate corollary of the following proposition on distributional communication complexity.

PROPOSITION 3.7. *For any distribution μ,*

$$2^{D_\varepsilon^\mu(F)} \geq \frac{1-2\varepsilon}{\operatorname{disc}_\mu(F)}. \tag{3.2.2}$$

Proof. Let π be any deterministic protocol that realizes $D_\varepsilon^\mu(F)$. Then π partitions the input space into $N \leq 2^{D_\varepsilon^\mu(F)}$ cylinder intersections, $\chi_1, \chi_2, \ldots, \chi_N$. Sum the error within each cylinder intersection,

$$\begin{aligned} \varepsilon &\geq \sum_{i=1}^{N} \frac{1}{2}\left(\mu(\chi_i) - \left|\sum_{x \in X_1 \times X_2 \times \cdots \times X_\ell} (-1)^{F(x)} \mu(x) \chi_i(x)\right|\right) \\ &\geq \frac{1}{2} - \frac{1}{2} N \operatorname{disc}_\mu(F). \end{aligned}$$

Rearranging the terms, we obtain (3.2.2). □

Combining Theorem 3.6 with the definition of $\mathsf{PP}(F)$ gives the following corollary.

Corollary 3.8. *Let $F\colon X_1 \times X_2 \times \cdots \times X_\ell \to \{0,1\}$ be a given communication problem. Then*

$$\mathsf{PP}(F) \geq \log \frac{2}{\operatorname{disc}(F)}.$$

In fact, discrepancy characterizes the communication complexity with weakly unbounded error PP fully due to Klauck [**76**],

$$\mathsf{PP}(F) \leq O\left(\log \frac{1}{\operatorname{disc}(F)} + \log n\right).$$

Let me pause and make a remark that the discrepancy method also gives a lower bound for the quantum communication complexity with or without prior entanglement, see [**83, 89**].

The *sign-rank* of a real matrix $A \in \mathbb{R}^{n\times m}$ with nonzero entries is the least rank of a matrix $B \in \mathbb{R}^{n\times m}$ such that $\operatorname{sgn} A_{i,j} = \operatorname{sgn} B_{i,j}$ for all i,j. In general, the sign-rank of a matrix can be vastly smaller than its rank. For example, consider the following nonsingular matrices of order $n \geq 3$:

$$\begin{bmatrix} 1 & & & & & \\ & 1 & & & \text{\Large 1} & \\ & & 1 & & & \\ & & & \ddots & & \\ \text{\Large -1} & & & & 1 & \\ & & & & & 1 \end{bmatrix}, \qquad \begin{bmatrix} 1 & & & & & \\ & 1 & & & \text{\Large -1} & \\ & & 1 & & & \\ & & & \ddots & & \\ \text{\Large -1} & & & & 1 & \\ & & & & & 1 \end{bmatrix}.$$

These matrices have sign-rank at most 2 and 3, respectively. Indeed, the first matrix has the same sign pattern as $[2(j-i)+1]_{i,j}$. The second has the same sign pattern as $[\langle v_i, v_j\rangle - (1-\varepsilon)]_{i,j}$, where $v_1, v_2, \ldots, v_n \in \mathbb{R}^2$ are arbitrary pairwise distinct unit vectors and ε is a suitably small positive real, cf. [**101**, Section 5]. As a matter of notational convenience, we extend the notion of sign-rank to Boolean functions $f\colon X \times Y \to \{0,1\}$ by defining $\operatorname{rk}_\pm(f) = \operatorname{rk}_\pm(M_f)$, where $M_f = [(-1)^{f(x,y)}]_{x\in X, y\in Y}$

is the matrix associated with f. A remarkable fact, due to Paturi and Simon [**101**], is that the sign-rank of a two-party communication problem fully characterizes its unbounded-error communication complexity.

THEOREM 3.9 (Paturi and Simon). *Let $F\colon X\times Y\to\{0,1\}$ be a given communication problem. Then*

$$\log \mathrm{rk}_{\pm}(F)\le \mathsf{UPP}(F)\le \log \mathrm{rk}_{\pm}(F)+2.$$

The first inequality is essentially a reformulation of Proposition 3.4. The second inequality follows similar ideas as expressing the sign matrix as a sum of rank one matrices $u^T\cdot v$ by rank decomposition and then interpreting $u^T(x), v(y)$ as the expectations of sending -1 or 1 (under suitable normalization).

As Corollary 3.8 and Theorem 3.9 show, the study of communication with unbounded and weakly unbounded error is in essence the study of discrepancy and sign-rank.

3.2.3. The pattern matrix method.

Discrepancy and sign-rank are difficult to analyze from first principles. The *pattern matrix method*, developed in [**114, 116**], is a technique that transforms lower bounds for polynomial approximation into bounds on discrepancy, sign-rank, and various other quantities in communication complexity. For our discrepancy bounds, we use the following special case of the pattern matrix method [**123**, Theorem 5.7 and equation (119)].

THEOREM 3.10 (Sherstov). *Let $f\colon\{0,1\}^n\to\{0,1\}$ be given. Consider the ℓ-party communication problem $F\colon(\{0,1\}^{nm})^\ell\to\{0,1\}$ given by $F=f\circ \mathrm{NOR}_m\circ \mathrm{AND}_\ell$. Then*

$$\mathrm{disc}(F)\le\left(\frac{c2^\ell\ell}{\sqrt{m}}\right)^{\deg_\pm(f)/2},$$

where $c > 0$ is a constant independent of n, m, ℓ, f.

We note that the case $\ell = 2$ of Theorem 3.10 is vastly easier to prove than the general statement; this two-party result can be found in [**125**, Theorem 7.3 and equation (7.3)]. For our sign-rank lower bounds, we use the following theorem implicit in [**117**].

THEOREM 3.11 (Sherstov, implicit). *Let $f\colon \{0,1\}^n \to \{0,1\}$ be given. Suppose that $\deg_\pm(f,\gamma) \geq d$, where γ and d are positive reals. Fix an integer $m \geq 2$ and define $F\colon \{0,1\}^{mn} \times \{0,1\}^{mn} \to \{0,1\}$ by $F = f \circ \mathrm{OR}_m \circ \mathrm{AND}_2$. Then*

$$\mathrm{rk}_\pm(F) \geq \gamma \left\lfloor \frac{m}{2} \right\rfloor^{d/2}.$$

For the reader's convenience, we give a detailed proof of Theorem 3.11 in Section 5.7.

CHAPTER 4

Communication against noise

In this chapter, we discuss our contribution to communication in noisy environments. We will study how to encode a protocol so that the encoded protocol can tolerate a maximum fraction of adversarial corruptions of insertions, deletions and substitutions. This is the interactive analogue of error correcting codes with optimal corruption rate.

4.1. Introduction

Consider the classical problem of transmitting a message over an unreliable channel. In its most general formulation, the problem features an omniscient and computationally unbounded adversary who controls the communication channel and can alter a small constant fraction of symbols that pass through the channel. The choice of symbols to corrupt is up to the adversary; the only guarantee is an a priori bound on the fraction of altered symbols, called the *corruption rate*. The sender's objective is to encode the message using a somewhat longer string so as to always allow the receiver to recover the original message. Shannon's problem is the subject matter of coding theory and has been extensively studied. In particular, for any constant $\varepsilon > 0$, it is known [**74**] how to encode an n-bit message using a string of $O(n)$ symbols from a constant-size alphabet such that the receiving party will recover the original message whenever the fraction of corrupted symbols is at most $\frac{1}{2} - \varepsilon$. In seminal work, Schulman [**109, 110, 111**] considered a generalization of Shannon's problem to the interactive setting. Here, two parties Alice and Bob communicate back and forth according to a communication protocol agreed upon in advance. Alice and Bob privately hold inputs X and Y, respectively, which dictate their behavior throughout the communication protocol. As before, the communication channel is controlled by an adversary who can change a small constant fraction of symbols as they transit through the channel. The goal is to overcome these corruptions by cleverly simulating the original protocol with some redundant communication, as follows. The simulation leaves Alice and Bob with a record of symbols exchanged between them, where Alice's record will generally disagree with Bob's due to interference by the adversary. Nevertheless, they each need to be able to determine, with no further communication, the sequence of symbols that would have been exchanged in the *original* protocol on the inputs X and Y in question. Ideally, Alice and Bob's simulation should use an alphabet of constant size and have communication cost within a constant factor of the original protocol.

A naïve solution to Schulman's problem is for Alice and Bob to encode their individual messages with an error-correcting code developed for Shannon's setting. This approach fails spectacularly because the adversary is only restricted by the total number of corruptions rather than the number of corruptions on a per-message basis. In particular, the adversary may choose a specific message from Alice to Bob and corrupt all symbols in it. As a result, the naïve solution cannot tolerate any corruption rate beyond $\frac{1}{m}$, where m is the total number of messages. Remarkably, Schulman [**111**] was able to show how to simulate any communication protocol with corruption rate up to $\frac{1}{240}$, using a constant-size alphabet and a constant-factor overhead in communication. Interactive coding has since evolved into a highly active research area with a vast literature on virtually every aspect of the problem, e.g., [**98, 26, 61, 20, 55, 80, 63, 24, 22, 62, 71, 51, 59, 6**], from corruption rate to communication overhead to computational complexity. We refer the reader to

Gelles [**58**] for an up-to-date survey. Of particular interest to us is the work of Braverman and Rao [**26**], who proved that any communication protocol can be simulated in Schulman's model with corruption rate up to $\frac{1}{4} - \varepsilon$ for any $\varepsilon > 0$, and established a matching impossibility result for corruption rate $\frac{1}{4}$. Analogous to Schulman [**111**], the simulation due to Braverman and Rao [**26**] uses a constant-size alphabet and increases the communication cost only by a constant factor.

In the canonical model discussed above, the adversary manipulates the communication channel by altering symbols. This type of manipulation is called a *substitution.* In a recent paper, Braverman, Gelles, Mao, and Ostrovsky [**25**] proposed a far-reaching generalization of the canonical model, whereby the adversary can additionally manipulate the channel by inserting and deleting symbols. As Braverman et al. point out, insertions and deletions are considerably more difficult to handle than substitutions even in the one-way setting of coding theory. To borrow their example, Schulman and Zuckerman's polynomial-time coding and decoding algorithms [**112**] for insertion and deletion errors can tolerate a corruption rate of roughly $\frac{1}{100}$, in contrast to the corruption rate of $\frac{1}{2} - \varepsilon$ or $\frac{1}{4} - \varepsilon$ (depending on the alphabet size) achievable in the setting of substitution errors alone [**74**]. As their main result, Braverman et al. [**25**] prove that any communication protocol can be simulated in the generalized model with substitutions, insertions, and deletions as along as the corruption rate does not exceed $\frac{1}{18} - \varepsilon$, for an arbitrarily small constant $\varepsilon > 0$. Analogous to previous work, the simulation of Braverman et al. uses a constant-size alphabet and increases the communication cost only by a multiplicative constant.

Braverman et al. [**25**] and Gelles [**58**] posed the problem of determining the highest possible corruption rate that can be tolerated in the generalized model, and of achieving that optimal rate for every protocol. We give a detailed solution to this problem, showing that any protocol can be simulated with corruption rate up to $\frac{1}{4} - \varepsilon$

for any $\varepsilon > 0$. Recall that this corruption tolerance is optimal even in the setting of substitutions alone.

4.1.1. The indels model. Following previous work, we focus on communication protocols in *canonical form.* In such a protocol, the communication proceeds in *rounds.* The number of rounds is the same on all inputs, and each round involves Alice sending a single symbol to Bob and Bob sending a symbol back to Alice. The canonical form assumption is without loss of generality since any protocol can be brought into canonical form at the expense of doubling its communication cost.

We now describe the model of Braverman et al. [**25**] in more detail. Naïvely, one may be tempted to give the adversary the power to delete or insert any symbol at any time. A moment's thought reveals that such power rules out any meaningful computation. Indeed, deleting a single symbol en route from Alice to Bob will stall the communication, forcing both parties to wait on each other indefinitely to send the next symbol. Conversely, inserting a symbol into the communication channel may result in crosstalk, with both parties trying to send a symbol at the same time. Braverman et al. [**25**] proposed a natural and elegant formalism, to which we refer as the *BGMO model*, that avoids these abnormalities. In their model, deletions and insertions occur in pairs, with every deletion immediately followed by an insertion. In other words, the BGMO model gives the adversary the capability to intercept any symbol σ in transit from one party to the other and insert a spurious symbol σ' in its place. Crucially, the adversary is free to decide which party will receive the inserted symbol. This makes it possible for the adversary to carry out two types of attacks, illustrated in Figure 4.1.1. In a *substitution attack*, the inserted symbol is routed the same way as the original symbol. Such an attack is precisely equivalent to a substitution in Schulman's model [**111**]. In an *out-of-sync attack*, on the other hand, the inserted symbol is delivered to the sender of the original symbol. From

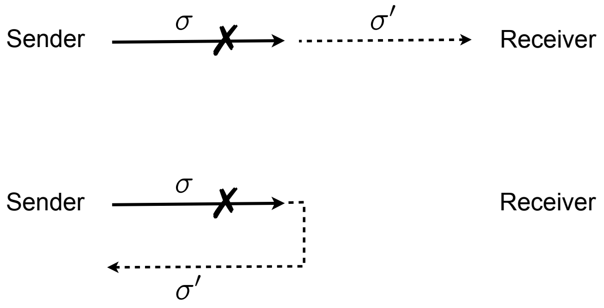

FIGURE 4.1.1. A substitution attack (top) and an out-of-sync attack (bottom).

the sender's point of view, an out-of-sync attack looks like a response from the other party, whereas that other party does not even know that any communication has taken place and continues to wait for an incoming symbol. Braverman et al. [**25**] examine a variety of candidate models, including some that are clock-driven rather than message-driven, and demonstrate that the BGMO model is essentially the only reasonable interactive formalism that allows deletions and insertions. It is important to note here that even though deletions and insertions in the BGMO model occur in pairs, the corruption pattern experienced by any given party can be an arbitrary sequence of deletions and insertions.

4.1.2. Our results.

For the purposes of defining the corruption rate, a deletion-insertion pair in the BGMO model counts as a single corruption. This means that with corruption rate δ, the adversary is free to carry out as many as δM attacks, where M is the worst-case number of sent symbols. The main result of our work is the following theorem, where $|\pi|$ denotes the worst-case communication cost of a protocol π.

Theorem 4.1. *Fix an arbitrary constant $\varepsilon > 0$, and let π be an arbitrary protocol with alphabet Σ. Then there exists a simulation for π with alphabet size $O(1)$ and communication cost $O(|\pi| \log |\Sigma|)$ that tolerates corruption rate $\frac{1}{4} - \varepsilon$ in the BGMO model.*

Theorem 4.1 matches an upper bound of $\frac{1}{4}$ on the highest possible corruption rate, due to Braverman and Rao [**26**], which holds even if the adversary is restricted to substitution attacks.

Theorem 4.1 is particularly generous in that it gives the adversary a flat budget of δM attacks, where δ is the corruption rate and M is the *maximum* number of sent symbols over all executions. Due to out-of-sync attacks, the number of symbols sent in a given execution may be substantially smaller than M. This can happen, for example, if the adversary uses out-of-sync attacks to force one of the parties to exit before his or her counterpart has reached the end of the simulation. In such case, the actual ratio of the number of attacks to the number of sent symbols may substantially exceed δ. This leads us to consider the following alternate formalism: with *normalized corruption rate* $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$, the number of substitution attacks and out-of-sync attacks in any given execution must not exceed an $\varepsilon_{\text{subs}}$ and ε_{oos} fraction, respectively, of the number of symbols sent in that execution. In this setting, we prove:

Theorem 4.2 (Normalized corruption rate). *Fix an arbitrary constant $\varepsilon > 0$, and let π be an arbitrary protocol with alphabet Σ. Then there exists a simulation for π with alphabet size $O(1)$ and communication cost $O(|\pi| \log |\Sigma|)$ that tolerates any normalized corruption rate $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$ in the BGMO model with*

$$\varepsilon_{\text{subs}} + \frac{3}{4}\varepsilon_{\text{oos}} \leq \frac{1}{4} - \varepsilon.$$

We show that Theorem 4.2, too, is optimal with respect to the normalized corruption rates that it tolerates (Section 4.5.9). In the interesting special case when the adversary is restricted to out-of-sync attacks, Theorem 4.2 tolerates normalized corruption rate $\frac{1}{3} - \varepsilon$ for any $\varepsilon > 0$. This contrasts with the maximum possible corruption rate that can be tolerated with substitutions alone, namely, $\frac{1}{4} - \varepsilon$. Thus, there is a precise technical sense in which substitution attacks are more powerful than out-of-sync attacks. As we will discuss shortly, however, the mere presence of out-of-sync attacks greatly complicates the analysis and requires a fundamentally different approach.

In Theorems 4.1 and 4.2, each player computes the transcript of the simulated protocol based on his or her *entire* record of sent and received symbols, from the beginning of time until the communication stops. In Section 4.5.8, we adapt Theorem 4.1 to the setting where Alice and Bob wish to know the answer by a certain round, according to each player's own counting. In particular, Braverman et al. [**25**] required each player to know the answer by round $(1 - 2\delta)N$, where N is the maximum number of rounds and δ is the corruption rate. With that requirement, we give a simulation that tolerates corruption rate $\frac{1}{6} - \varepsilon$ for any $\varepsilon > 0$, which is optimal by the impossibility result in [**25**, Theorem G.1].

4.1.3. Background on interactive coding. In what follows, we review relevant previous work [**111, 26, 25**] on interactive coding and contrast it with our approach. A key tool in this line of research is a *tree code*, a coding-theoretic primitive developed by Schulman [**111**]. Let Σ_{in} and Σ_{out} be nonempty finite alphabets. A tree code is any length-preserving map $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ with the property that for any input string $s \in \Sigma_{\text{in}}^*$ and any $i = 1, 2, 3, \ldots$, the first i symbols of the codeword $C(s)$ are completely determined by the first i symbols of the input string s. A tree code has a natural representation as an infinite tree in which every vertex has arity $|\Sigma_{\text{in}}|$ and every edge is labeled with a symbol from Σ_{out}. To compute the codeword corresponding to a given input string $s = s_1 s_2 \dots s_k$, one starts at the root and walks down the tree for k steps, choosing at the i^{th} step the branch that corresponds to s_i. The sought codeword $C(s)$, then, is the concatenation of the edge labels along this path. Tree codes are well-suited for encoding interactive communication because Alice and Bob must compute and send symbols one at a time, based on each other's responses, rather than all at once at the beginning of the protocol. In more detail, if Alice has used a tree code C to send Bob $s_1, s_2, \dots, s_{k-1}$ and now wishes to send him s_k, she need only send the k^{th} symbol of $C(s_1 s_2 \dots s_k)$ rather than all of $C(s_1 s_2 \dots s_k)$. This works because by the defining properties of a tree code, the first $k-1$ symbols of $C(s_1 s_2 \dots s_k)$ are precisely $C(s_1 s_2 \dots s_{k-1})$ and are therefore known to Bob already. To additionally cope with adversarial substitutions, Schulman used tree codes in which different codewords are "far apart." More precisely, for any two input strings $s, s' \in \Sigma_{\text{in}}^*$ of equal length with $s_1 s_2 \dots s_k = s'_1 s'_2 \dots s'_k$ but $s_{k+1} \neq s'_{k+1}$, the codewords $C(s)$ and $C(s')$ disagree in a $1-\alpha$ fraction of positions beyond the k^{th}. Schulman [**111**] showed the existence of such tree codes for any $\alpha > 0$, where the size of the output alphabet depends only on α and the input alphabet. Figure 4.1.2 (left) offers an illustration of the distance property for tree codes: the concatenation of the labels on the solid path should disagree with the concatenation of the labels on the dashed path in a $1 - \alpha$ fraction of positions. Finally, when attempting to recover the codeword from a corrupted string $y \in \Sigma_{\text{out}}^*$, one outputs the codeword of length $|y|$ that is closest to y in Hamming distance. This recovery procedure produces the true codeword whenever y is sufficiently close to some codeword in *suffix distance*, a distance on strings that arises in a natural way from tree code properties.

We now review protocol terminology. Fix a deterministic protocol π in canonical form that Alice and Bob need to simulate on their corresponding inputs X and Y. Let Σ and n denote the alphabet and the communication cost of π, respectively.

Associated to π is a tree of depth n called the *protocol tree for* π. Each vertex in this tree corresponds to the state of the protocol at some point in time, with the root corresponding to the initial state before any symbols have been exchanged, and each leaf corresponding to a final state when the communication has ended. Each internal vertex has arity $|\Sigma|$, corresponding to all possible symbols that can be transmitted at that point. Execution of π corresponds to a walk down the protocol tree, as follows. A given input X for Alice makes available precisely one outgoing edge for every internal vertex of even depth, corresponding to the symbol that she would send if the execution were to arrive at that vertex. Similarly, an input Y for Bob makes available precisely one outgoing edge for every internal vertex of odd depth. To execute π, Alice and Bob walk down the protocol tree one edge at a time, at each step selecting the edge that is dictated by the input of the player whose turn it is to speak. In this chapter, we will assume that output of the protocol π on a given pair of inputs X, Y, is the complete sequence of symbols exchanged between Alice and Bob on that pair of inputs. We make this somewhat usual assumption here because in this chapter our focus is to encode a protocol instead of computing a function.

We emphasize that there is no relation whatsoever between protocol trees and trees representing tree codes. They are structurally unrelated and play entirely different roles in the simulation of a protocol over an unreliable channel.

Given protocols π and Π with input space $\mathcal{X} \times \mathcal{Y}$, we say that Π *simulates* π if $\pi(X, Y) = f(\Pi(X, Y))$ for some fixed function f and all inputs $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$. To illustrate, any protocol π with alphabet Σ can be simulated in the natural manner by a protocol Π with the binary alphabet $\{0, 1\}$ and communication cost $|\Pi| \leq |\pi| \max\{1, \lceil \log |\Sigma| \rceil\}$. Observe that the "simulates" relation on protocols is transitive. A protocol π is said to be *in canonical form* if the following two conditions hold: (i) the number of symbols exchanged between Alice and Bob is an even integer and

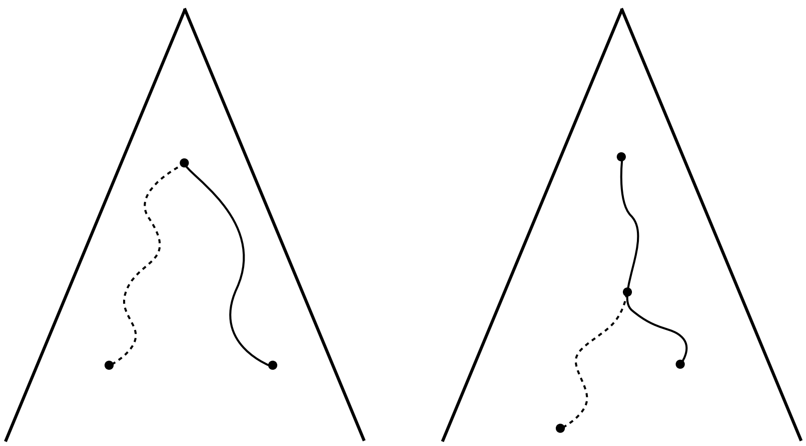

Figure 4.1.2. Distance constraints for codewords in a tree code (left) and an edit distance tree code (right).

is the same for all inputs $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$; (ii) Alice and Bob take turns sending each other one symbol at a time, with Alice sending the first symbol. A moment's thought reveals that any protocol π can be simulated by a protocol in canonical form with the same alphabet and at most double the communication cost.

In this chapter, we view the transfer of an alphabet symbol from one party to the other as an atomic operation to which we refer as a *transmission*. We also intentionally avoid the term "message" in this chapter because it is ambiguous as to the length of the content.

4.1.4. The Braverman–Rao simulation. We are now in a position to describe the simulation of Braverman and Rao [**26**] for the model with adversarial substitutions. Using the tree view of communication, we can identify Alice's input X with a set E_X of outgoing edges for the protocol tree vertices at even depths, one such edge per vertex. Analogously, Bob's input Y corresponds to a set E_Y of outgoing edges for the vertices at odd depths. Execution of π, then, corresponds to identifying the unique root-to-leaf path made up of edges in $E_X \cup E_Y$. In Braverman and Rao's simulation, all communication is encoded and decoded using a tree code with the parameter $\alpha > 0$ set to a small constant. The simulation amounts to Alice and Bob taking turns sending each other edges from their respective sets E_X and E_Y. When it is Alice's turn to speak, she decodes the edge sequence received so far and attempts to extend the path made up of her sent and received edges by another edge from E_X, communicating this new edge to Bob. Bob acts analogously. When the communication stops, Alice decodes her complete sequence of received edges, identifies the first prefix of that sequence whose edges along with E_X contain a root-to-leaf path, and takes this root-to-leaf path to be the transcript of π on the given pair of inputs. Bob, again, acts analogously.

In the described simulation, the edge that a player sends at any given point may be irrelevant but it is never incorrect. In particular, Alice and Bob make progress in every round where they correctly decode the edge sequences that they have received so far. Braverman and Rao use a relation between suffix distance and Hamming distance to argue that with overall corruption rate $\frac{1}{4} - \varepsilon$, Alice decodes her received edge sequence correctly more often than half of the time, and likewise for Bob. This means that there are a considerable number of rounds where Alice and Bob *both* decode their received sequences correctly. It follows that at some point t^*, Alice and Bob will have exchanged every edge in the root-to-leaf path in $E_X \cup E_Y$. As a final ingredient, the authors of [**26**] argue that the adversary's remaining budget for corruptions beyond time t^* cannot "undo" this progress, in the sense that at the end of the communication Alice and Bob will correctly decode a prefix that contains the root-to-leaf path in $E_X \cup E_Y$.

4.1.5. The BGMO simulation. We now describe the simulation of Braverman et al. [**25**] in the BGMO model with substitutions, insertions, and deletions. The authors of [**25**] draw inspiration from the classic work of Levenshtein [**87**], who developed codes that allow recovery from insertions and deletions in the noninteractive setting. Recall that when coding for substitution errors, one uses codewords that are far apart in Hamming distance [**74**]. Analogously, Levenshtein used codewords that are far apart in *edit distance*, defined for a pair of strings as the minimum number of insertions and deletions needed to transform one string into the other. To handle interactive communication, then, it is natural to start as Braverman et al. do with a tree code in which the codewords are far apart in edit distance rather than Hamming distance. They authors of [**25**] discover, however, that it is no longer sufficient to have distance constraints for pairs of codewords of the *same* length. Instead, for any two paths of arbitrary lengths that cross to form a lambda shape, such as the solid and dashed paths in Figure 4.1.2 (right), the associated codeword segments need to be far apart in edit distance. Braverman et al. establish the existence of such *edit distance tree codes* and develop a notion of suffix distance for them, thus providing a sufficient criterion for the recovery of the codeword from a corrupted string.

Algorithmically, the BGMO simulation departs from Braverman and Rao's in two ways. First, all communication is encoded and decoded using an edit distance tree code. Second, a different mechanism is used to decide which leaf of the protocol tree for π to output, whereby each player keeps a tally of the number of times any given leaf has been reached during the simulation and outputs the leaf with the highest tally. The resulting analysis is quite different from [**26**], out-of-sync attacks being the main source of difficulty. Braverman et al. start by showing that each player correctly decodes his or her received sequence of edges often enough over the course of the simulation. This does not imply progress, however. Indeed, all of Alice's correct decodings may conceivably precede all of Bob's, whereas progress is only guaranteed when the players' correct decodings are interleaved. To prove that this interleaving takes place, Braverman et al. split the simulation into n progress intervals, corresponding to the length of the longest segment recovered so far from the root-to-leaf path in $E_X \cup E_Y$. They use an amortized analysis to argue that the number of unsuccessful decodings per interval is small on the average, allowing Alice and Bob to reach the leaf on the root-to-leaf path in $E_X \cup E_Y$ at some point in the simulation. They finish the proof by arguing that the players subsequently revisit this leaf often enough that its tally outweighs that of any other leaf.

4.1.6. Our approach. There are several obstacles to improving the corruption tolerance from $\frac{1}{18} - \varepsilon$ in Braverman et al. [**25**] to an optimal $\frac{1}{4} - \varepsilon$. Some of these obstacles are of a technical nature, whereas others require a fundamental shift in approach and analysis. In the former category, we develop edit distance tree codes with stronger guarantees. Specifically, Braverman et al. use tree codes with the property that for any two paths that cross to form a lambda shape in the code tree, the edit distance between the associated codeword segments is at least a $1 - \alpha$ fraction of the length of the *longer* path. We prove the existence of tree codes that guarantee a stronger lower bound on the edit distance, namely, a $1 - \alpha$ fraction of the *sum* of the lengths of the paths. This makes intuitive sense because the typical edit distance between randomly chosen strings of lengths ℓ_1 and ℓ_2 over a nontrivial alphabet is approximately $\ell_1 + \ell_2$ rather than $\max\{\ell_1, \ell_2\}$; cf. Proposition 4.4. Our second improvement concerns the decoding process. The notion of suffix distance used by Braverman et al. is not flexible enough to support partial recovery of a codeword. We define a more general notion that we call *k-suffix distance* and use it to give a sufficient criterion for the recovery of the first k symbols of the codeword from a corrupted string. This makes it possible to replace the tally-based output criterion of Braverman et al. with a more efficient mechanism, whereby Alice and Bob compute their output based on a *prefix* on the received edge sequence rather than the entire sequence.

The above technical improvements fall short of achieving an optimal corruption rate of $\frac{1}{4} - \varepsilon$. The fundamental stumbling block is the presence of out-of-sync attacks. For one thing, Alice and Bob's transmissions can now be interleaved in a complex way, and the basic notion of a round of communication is no longer available. Out-of-sync attacks also break the symmetry between the two players in that it is now possible for one of them to receive substantially fewer symbols than the other. Finally, by directing a large number of out-of-sync attacks at one of the players, the adversary can force the simulation to stop early and thereby increase the effective error rate well beyond $\frac{1}{4} - \varepsilon$. These are good reasons to doubt the existence of a simulation that tolerates corruption rate $\frac{1}{4} - \varepsilon$ with substitutions, insertions, and deletions.

Our approach is nevertheless based on the intuition that out-of-sync attacks should actually *help* the analysis because they spread the brunt of a corruption between the two players rather than heaping it all on a single player. Indeed, the deletion that results from an out-of-sync attack only affects the receiver, whereas the insertion only affects the sender. This contrasts with substitution attacks, where the deletions and insertions affect exclusively the receiver. With this in mind, convexity considerations suggest that out-of-sync attacks may actually be less damaging overall than substitution attacks. To bear out this intuition, we introduce a "virtual" view of communication that centers around the *events* experienced by Alice and Bob (namely, insertions, deletions, and successful deliveries) rather than the *symbols* that they send. In this virtual view, the length of a time interval and the associated error rate are defined in terms of the number of alternations in events rather than in terms of the number of sent symbols. Among other things, the virtual view restores the symmetry between Alice and Bob and makes it impossible for the adversary to shorten the simulation using out-of-sync attacks. By way of analysis, we start by proving that corruption rate $\frac{1}{4} - \varepsilon$ translates into virtual corruption rate $\frac{1}{4} - \Omega(\varepsilon)$. Next, we split the simulation into n progress intervals, corresponding to the length of the longest segment recovered so far from the root-to-leaf path in $E_X \cup E_Y$, and a final interval that encompasses the remainder of the simulation. We bound the virtual length of each interval in terms of the number of corruptions and successful decodings. We then contrast this bound with the virtual length of the overall simulation, which unlike actual length is never smaller than the simulation's worst-case communication complexity. Using the previously obtained $\frac{1}{4} - \Omega(\varepsilon)$ upper bound on the virtual corruption rate, we argue that Alice and Bob successfully output the root-to-leaf path in $E_X \cup E_Y$ when their communication stops.

4.2. Preliminaries

We start with a review of the technical preliminaries.

4.2.1. Edit distance. Recall that the asterisk $*$ is a reserved symbol that does not appear in any alphabet Σ in this manuscript. For a string $v \in (\Sigma \cup \{*\})^*$, we let $*(v)$ and $\overline{*}(v)$ denote the number of asterisks and non-asterisk symbols in v, respectively:

$$\begin{aligned} *(v) &= |\{i : v_i = *\}|, \\ \overline{*}(v) &= |\{i : v_i \neq *\}|. \end{aligned}$$

In particular, $*(v) + \overline{*}(v) = |v|$. We let $\not*(v)$ stand for the string of length $\overline{*}(v)$ obtained from v by deleting the asterisks. For example, $\not*(*ab*aa) = abaa$ and $\not*(*) = \varepsilon$ for any alphabet symbols a, b.

An *alignment* for a given pair of strings $s, r \in \Sigma^*$ is a pair of strings $S, R \in (\Sigma \cup \{*\})^*$ with the following properties:

$$\begin{aligned}
&|S| = |R|, && \\
&\not\ast(S) = s, && \\
&\not\ast(R) = r, && \\
&R_i \neq * \ \vee \ S_i \neq * && (i = 1, 2, 3, \ldots, |S|), \\
&(R_i \neq * \ \wedge \ S_i \neq *) \quad \Longrightarrow \quad R_i = S_i && (i = 1, 2, 3, \ldots, |S|).
\end{aligned}$$

To better distinguish alignments from ordinary strings, we reserve uppercase symbols for the former and lowercase for the latter. We write $S \parallel R$ to indicate that S and R are an alignment for some pair of strings. For an alignment $S \parallel R$, the strings $S|_A, R|_A$ for any given subset A of indices also form an alignment, to which we refer as a *subalignment* of $S \parallel R$.

The notion of a string alignment arises in an auxiliary capacity in the context of edit distance. Specifically, the *edit distance* between strings $s, r \in \Sigma^*$ is denoted $\mathrm{ED}(s, r)$ and is given by

$$\mathrm{ED}(s, r) = \min_{S \parallel R} \{*(S) + *(R)\},$$

where the minimum is over all alignments for s, r. Letting $\mathrm{LCS}(s, r)$ denote the length of the longest common subsequence of s and r, we immediately have

$$\mathrm{ED}(s, r) = |s| + |r| - 2\,\mathrm{LCS}(s, r). \tag{4.2.1}$$

The following equivalent definition is frequently useful: $\mathrm{ED}(s, r)$ is the minimum number of insertion and deletion operations necessary to transform s into r. In this equivalence, an alignment $S \parallel R$ represents a specific way to transform s into r, indicating the positions of the insertions ($S_i = *, R_i \neq *$), deletions ($S_i \neq *, R_i = *$), and unchanged symbols ($S_i = R_i \neq *$). The operational view of edit distance shows that it is a metric, with all strings s, r, t obeying

$$\mathrm{ED}(s, r) = \mathrm{ED}(r, s), \tag{4.2.2}$$

$$\mathrm{ED}(s, r) + \mathrm{ED}(r, t) \leq \mathrm{ED}(s, t). \tag{4.2.3}$$

Another property of edit distance is as follows.

Proposition 4.3. *For any strings $u, v \in \Sigma^*$,*

$$\mathrm{ED}(u, v) \geq ||u| - |v||.$$

In particular,

$$\mathrm{ED}(u, v) = ||u| - |v||$$

whenever u is a subsequence of v or vice versa.

Proof. The proposition is immediate from (4.2.1). An alternate approach is to appeal to the operational view of edit distance, as follows. An insertion or deletion changes the length of a string by at most 1. Therefore, at least $\max\{|u| - |v|, |v| - |u|\} = ||u| - |v||$ operations are needed to transform u into v. If one of the strings is a *subsequence* of the other, then either of them can clearly be transformed into the other using $||u| - |v||$ deletions or $||u| - |v||$ insertions. □

By definition, the edit distance between a pair of strings of lengths n and m is at most $n + m$. We now show that this trivial upper bound is essentially tight when the strings are chosen uniformly at random over an alphabet of nonnegligible size.

Proposition 4.4. *For any nonnegative integers n and m and any $0 < \alpha \le 1$,*

$$\mathop{\mathbf{P}}_{\substack{u\in\Sigma^n\\ v\in\Sigma^m}}[\mathrm{ED}(u,v) \le (1-\alpha)(n+m)] \le \left(\frac{e}{\alpha\sqrt{|\Sigma|}}\right)^{\alpha(n+m)}.$$

Proof. We may assume that

$$\frac{e}{\alpha\sqrt{|\Sigma|}} \le 1, \tag{4.2.4}$$

the proposition being trivial otherwise. Letting $\ell = \lceil \alpha(n+m)/2 \rceil$, we have

$$\begin{aligned}
\mathop{\mathbf{P}}_{\substack{u\in\Sigma^n\\ v\in\Sigma^m}}[\mathrm{ED}(u,v) \le (1-\alpha)(n+m)] &= \mathop{\mathbf{P}}_{\substack{u\in\Sigma^n\\ v\in\Sigma^m}}[\mathrm{LCS}(u,v) \ge \ell] \\
&\le \binom{n}{\ell}\binom{m}{\ell} \cdot \frac{|\Sigma|^{\ell} \cdot |\Sigma|^{n-\ell} \cdot |\Sigma|^{m-\ell}}{|\Sigma|^{n+m}} \\
&\le \binom{n+m}{2\ell} \cdot \frac{1}{|\Sigma|^{\ell}} \\
&\le \left(\frac{e\,(n+m)}{2\ell} \cdot \frac{1}{\sqrt{|\Sigma|}}\right)^{2\ell} \\
&\le \left(\frac{e}{\alpha\sqrt{|\Sigma|}}\right)^{2\lceil \alpha(n+m)/2 \rceil} \\
&\le \left(\frac{e}{\alpha\sqrt{|\Sigma|}}\right)^{\alpha(n+m)},
\end{aligned}$$

where the first and last steps follow from (4.2.1) and (4.2.4), respectively. □

4.2.2. Suffix distance. We now discuss several other measures of distance for alignments and strings. For an alignment $S \parallel R$, define

$$\Delta(S,R) = \frac{*(S) + *(R)}{\overline{*}(S)}.$$

This quantity ranges in $[0, \infty]$, with the extremal values taken on. For example, $\Delta(\varepsilon, \varepsilon) = \Delta(a, a) = 0$ and $\Delta(*, a) = \infty$, where a is any alphabet symbol. The definition of Δ is motivated in large part by its relation to edit distance:

FACT 4.5. *For any alignment $S \parallel R$ with $\Delta(S, R) < \infty$,*

$$\mathrm{ED}(\cancel{*}(S), \cancel{*}(R)) \le \Delta(S, R) \cdot \overline{*}(S).$$

Proof. Immediate from the definitions of ED and Δ. $\square$

The *suffix distance* for an alignment $S \parallel R$ is given by

$$\mathrm{SD}(S, R) = \max_{i \ge 1} \, \Delta(S_{\ge i}, R_{\ge i}).$$

This notion was introduced recently by Braverman et al. [**25**], inspired in turn by an earlier notion of suffix distance due to Schulman [**111**]. In our work, we must consider a more general quantity yet. Specifically, we define $\mathrm{SD}_k(S, R)$ for $0 \le k \le \infty$ to be the maximum $\Delta(S_{\ge i}, R_{\ge i})$ over all indices i for which $\overline{*}(S_{<i}) < k$, with the convention that $\mathrm{SD}_k(S, R) = 0$ for $k = 0$. As functions, we have

$$0 = \mathrm{SD}_0 \le \mathrm{SD}_1 \le \mathrm{SD}_2 \le \mathrm{SD}_3 \le \cdots \le \mathrm{SD}_\infty = \mathrm{SD}\,. \tag{4.2.5}$$

We generalize the above definitions to strings $s, r \in \Sigma^*$ by letting

$$\mathrm{SD}(s, r) = \min_{S \parallel R} \, \mathrm{SD}(S, R), \tag{4.2.6}$$

$$\mathrm{SD}_k(s, r) = \min_{S \parallel R} \, \mathrm{SD}_k(S, R), \tag{4.2.7}$$

where in both cases the minimum is over all alignments $S \parallel R$ for s, r. Since there are only finitely many alignments for any pair of strings s and r, the quantities (4.2.6) and (4.2.7) can be computed in finite time.

4.2.3. Trees and tree codes. In a given tree, a *rooted path* is any path that starts at the root of the tree. The *predecessors* of a vertex v are any of the vertices on the path from the root to v, including v itself. We analogously define the *predecessors* of an edge e to be any of the edges of the rooted path that ends with e, including e itself. A *proper predecessor* of a vertex v is any predecessor of v other than v itself; analogously for edges. In keeping with standard practice, we draw trees with the root at the top and the leaves at the bottom. Accordingly, we define the *depth* of a vertex v as the length of the path from the root to v. Similarly, the *depth* of an edge e is the length of the rooted path that ends with e. We say that a given vertex v is *deeper* than another vertex u if the depth of v is larger than the depth of u; and likewise for edges.

Fix alphabets Σ_{in} and Σ_{out}. A *tree code* is any length-preserving map $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ such that the first i symbols of the output are completely determined by the first i symbols of the input. Formally,

$$\begin{aligned}
&|C(x)| = |x|,\\
&(C(x))_{\leq i} = C(x_{\leq i}), \qquad\qquad i = 0, 1, 2, \ldots,
\end{aligned}$$

for all $x \in \Sigma_{\text{in}}^*$. Recall that the *codewords* of C are the elements of $C(\Sigma_{\text{in}}^*)$, i.e., the strings $y \in \Sigma_{\text{out}}^*$ such that $y = C(x)$ for some x. A tree code can be represented as an infinite rooted tree in which each node has precisely $|\Sigma_{\text{in}}|$ outgoing edges, and each edge is labeled with a symbol from Σ_{out}. To compute $C(x)$ for a given string $x \in \Sigma_{\text{in}}^*$, one starts at the root and walks down the tree for $|x|$ steps, taking the edge corresponding to x_i in the i^{th} step. Then $C(x)$ is the concatenation of the $|x|$ edge labels, in the order they were encountered during the walk. If there is an a priori bound n on the length of the input string, as in this manuscript, it is sufficient to work with the restriction of the tree code to strings of length up to n. We refer to such a restriction as a *tree code of depth n*.

To allow decoding in the presence of errors, structural properties of a tree code must ensure that the encodings of distinct strings are sufficiently far apart. How this is formalized depends on the kinds of errors that must be tolerated. Previous work has considered substitution errors [**111, 26**] and more recently insertions and deletions [**25**]. We work in the latter setting and adopt structural constraints similar to those in [**25**].

DEFINITION 4.6 (α-violation). Fix a tree code $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ and a real $0 \le \alpha < 1$. A quadruple (A, B, D, E) of vertices in the tree representation of C form an *α-violation* if:

(i) B is the deepest common predecessor of D and E;

(ii) A is any predecessor of B; and

(iii) $\text{ED}(AD, BE) < (1-\alpha)(|AD|+|BE|)$, where $AD \in \Sigma_{\text{out}}^*$ is the concatenation of the code symbols along the path from A to D, and analogously $BE \in \Sigma_{\text{out}}^*$ is the concatenation of the code symbols along the path from B to E.

An *α-good code* is any tree code C for which no vertices A, B, D, E in its tree representation form an α-violation.

Definition 4.6 is illustrated in Figure 4.2.1. This definition strengthens an earlier formalism due to Braverman et al. [**25**], in which the inequality $\text{ED}(AD, BE) < (1-\alpha)\max\{|AD|, |BE|\}$ played the role of our constraint (iii). The strengthening is essential to the tight results of our work.

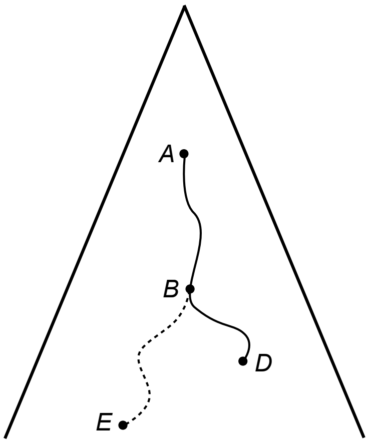

FIGURE 4.2.1. A quadruple of vertices A, B, D, E involved in an α-violation.

REMARK 4.7. Observe that A, B, D, E can form an α-violation for $0 \le \alpha < 1$ only when

$$D \neq E,$$
$$B \neq E.$$

Indeed, suppose that one or both of these conditions fail. Then $BE = \varepsilon$ and therefore

$$\begin{aligned} \mathrm{ED}(AD, BE) &= \mathrm{ED}(AD, \varepsilon) \\ &= |AD| \\ &= |AD| + |BE| \\ &\ge (1-\alpha)(|AD| + |BE|), \end{aligned}$$

where the second step follows from Proposition 4.3.

As the next observation shows, an α-good code allows for the unique decoding of every codeword.

FACT 4.8. *Let $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ be any α-good code, where $0 \leq \alpha < 1$. Then C is one-to-one.*

Proof. It will be convenient to prove the contrapositive. Let $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ be a tree code such that

$$C(x') = C(x''), \tag{4.2.8}$$

$$x' \neq x'' \tag{4.2.9}$$

for some strings $x', x'' \in \Sigma_{\text{in}}^*$. Let x be the longest common prefix of x' and x''. Consider the vertices B, D, E in the tree representation of C that correspond to the input strings $x, x', x'' \in \Sigma_{\text{in}}^*$, respectively. Then

$$\text{ED}(BD, BE) = 0 < (1-\alpha)(|BD| + |BE|),$$

where the first and second steps in the derivation follow from (4.2.8) and (4.2.9), respectively. Thus, the quadruple (B, B, D, E) forms an α-violation in C. □

The following theorem, proved using the probabilistic method, ensures the existence of α-good codes with good parameters.

THEOREM 4.9. *For any alphabet Σ_{in}, any $0 < \alpha < 1$, and any integer $n \geq 0$, there is an α-good code $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ of depth n with*

$$|\Sigma_{\text{out}}| = \left\lceil \frac{(10|\Sigma_{\text{in}}|)^{1/\alpha}\, e}{\alpha} \right\rceil^2.$$

This theorem and its proof are adaptations of an earlier result due to Braverman et al. [**25**]. For the reader's convenience, we provide a complete and self-contained proof of Theorem 4.9 in Section 4.3.2.

4.2.4. The corruption model. We adopt the corruption model introduced by Braverman et al. [**25**]. In this model, the communication channel between Alice and Bob is controlled by an omniscient and computationally unbounded adversary. In particular, the adversary knows Alice and Bob's protocol and their inputs. The adversary can interfere with a transmission in two different ways, illustrated in Figure 4.1.1. In a *substitution attack*, the adversary intercepts the sender's symbol σ and replaces it with a different symbol σ', which is then delivered to the receiver. In an *out-of-sync attack*, the adversary intercepts the sender's symbol σ, discards it, and then sends a spurious symbol σ' back to the sender in lieu of a response. Both a substitution attack and an out-of-sync attack involve the deletion of a symbol from the channel followed immediately by the insertion of a symbol; what makes these attacks different is how the inserted symbol is routed. On arrival, symbols manipulated by the adversary are indistinguishable from correct deliveries. As a result, Alice and Bob cannot in general tell on receipt of a transmission if it is corrupted. We remind the reader that a transmission is an atomic operation from the standpoint of interference by the adversary: either a transmission is delivered correctly and in full, or else an attack takes place and the transmission is considered to be corrupted.

Execution of a protocol is now governed not only by Alice and Bob's inputs but also by the adversary's actions. Our objective is to faithfully simulate any protocol π with only a constant-factor increase in communication cost. Our simulations will all be in canonical form, with Alice and Bob taking turns sending one symbol at a time. There are two immediate benefits to this strict alternation. First, it guarantees that the adversary cannot force *crosstalk*, with Alice and Bob attempting to send a transmission at the same time. Second, canonical form guarantees that the adversary cannot cause Alice and Bob both to *stall*, i.e., wait indefinitely on each other to send the next message. In particular, canonical form ensures that at least one of the parties is able to run the protocol to completion. The adversary may still force one of the parties to stall, e.g., by carrying out an out-of-sync attack during the next-to-last transmission. We consider an execution of the protocol to be complete as soon as the communication has stopped, due to Alice or Bob (or both) terminating.

With the adversary present, we must revisit the notion of protocol output. We define the *output* of a player in a particular execution to be the complete sequence of symbols, ordered chronologically, that player sends and receives over the course of the execution. There is a minor technicality to address regarding which received symbols are counted toward a player's output. Due to out-of-sync attacks, Alice and Bob need not always be in agreement about how many rounds of communication they have completed. As a result, it may happen that one of the players expects the communication to continue when the other has already exited. In that case, the latter player may have one last symbol addressed to him which he or she will never retrieve from the communication channel. Since that symbol is not accessible to the player, we do not count it toward his or her input. With this minor clarification, we are prepared for formalize our notion of an interactive coding scheme.

Definition 4.10 (Coding scheme). Let π be a given protocol with input space $\mathcal{X}\times\mathcal{Y}$. We say that protocol Π is an *interactive coding scheme for π that tolerates corruption rate ε* if:

(i) Π has input space $\mathcal{X}\times\mathcal{Y}$ and is in canonical form;

(ii) when Π is executed on a given pair of inputs $(X,Y)\in\mathcal{X}\times\mathcal{Y}$, the adversary is allowed to subject any transmission in Π to a substitution attack or out-of-sync attack, up to a total of at most $\varepsilon|\Pi|$ attacks;

(iii) there exist functions f', f'' such that for any pair of inputs $(X,Y)\in\mathcal{X}\times\mathcal{Y}$ and any allowable behavior by the adversary, Alice's output a and Bob's output b satisfy $f'(a)=f''(b)=\pi(X,Y)$.

In this formalism, the functions f' and f'' allow Alice and Bob to interpret their respective outputs as an output of the simulated protocol π, with the requirement that these interpretations by Alice and Bob match the actual output of π on the corresponding pair of inputs.

The previous definition gives the adversary a budget of $\varepsilon|\Pi|$ attacks, where $|\Pi|$ is the *maximum* length of any execution of Π. This flat budget applies even to executions that are significantly shorter than $|\Pi|$, as may happen due to out-of-sync attacks. This motivates us to define a second model, where the number of attacks in any given execution is bounded by a fraction of the *actual* length of that execution.

DEFINITION 4.11 (Coding scheme with normalized corruption rate). Let π be a given protocol with input space $\mathcal{X} \times \mathcal{Y}$. We say that protocol Π is an *interactive coding scheme for π that tolerates normalized corruption rate* $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$ if:

(i) Π has input space $\mathcal{X} \times \mathcal{Y}$ and is in canonical form;

(ii) when Π is executed on a given pair of inputs $(X, Y) \in \mathcal{X} \times \mathcal{Y}$, the adversary is allowed to subject any transmission in Π to a substitution attack or out-of-sync attack, where

- the number of substitution attacks in any execution is at most an $\varepsilon_{\text{subs}}$ fraction of the total number of transmissions in that execution, and
- the number of out-of-sync attacks in any execution is at most an ε_{oos} fraction of the total number of transmissions in that execution;

(iii) there exist functions f', f'' such that for any pair of inputs $(X, Y) \in \mathcal{X} \times \mathcal{Y}$ and any allowable behavior by the adversary, Alice's output a and Bob's output b satisfy $f'(a) = f''(b) = \pi(X, Y)$.

In this chapter, we will obtain an interactive coding scheme that achieves optimal corruption tolerance in both models (Definition 4.10 and 4.11).

4.3. Auxiliary results

We now prove a number of technical results on suffix distance and tree codes that are used in the design and analysis of our interactive coding schemes. Some of these results are new and some are adapted from previous work [**111, 26, 25**].

4.3.1. Bounds for suffix distance. Here, we collect several lower and upper bounds on suffix distance. We start with a proposition that gives bounds for alignments in terms of their subalignments.

PROPOSITION 4.12. *Let $S' \parallel R'$ and $S'' \parallel R''$ be given alignments. Then:*

(i) $\Delta(S'S'', R'R'') \leq \max\{\Delta(S', R'), \Delta(S'', R'')\}$;

(ii) $\Delta(S'S'', R'R'') \geq \min\{\Delta(S', R'), \Delta(S'', R'')\}$;

(iii) $\mathrm{SD}(S'S'', R'R'') \leq \max\{\mathrm{SD}(S', R'), \mathrm{SD}(S'', R'')\}$;

(iv) $\mathrm{SD}_k(S'S'', R'R'') \leq \max\{\mathrm{SD}_k(S', R'), \Delta(S'', R'')\}$ *for* $k \leq \overline{*}(S')$.

Proof. (i), (ii) There are two cases to consider. If $\overline{*}(S') > 0$ and $\overline{*}(S'') > 0$, we have

$$\begin{aligned}
\Delta(S'S'', R'R'') &= \frac{*(S'S'') + *(R'R'')}{\overline{*}(S'S'')} \\
&= \frac{*(S') + *(R')}{\overline{*}(S') + \overline{*}(S'')} + \frac{*(S'') + *(R'')}{\overline{*}(S') + \overline{*}(S'')} \\
&= \frac{\overline{*}(S')}{\overline{*}(S') + \overline{*}(S'')} \cdot \Delta(S', R') + \frac{\overline{*}(S'')}{\overline{*}(S') + \overline{*}(S'')} \cdot \Delta(S'', R'').
\end{aligned}$$

In other words, $\Delta(S'S'', R'R'')$ is a weighted average of $\Delta(S', R')$ and $\Delta(S'', R'')$ and therefore lies between the minimum and maximum of these quantities.

For the complementary case, by symmetry we may assume that $\overline{*}(S') = 0$. If $S' = \varepsilon$, then $\Delta(S'S'', R'R'') = \Delta(S'', R'')$ and therefore (i) and (ii) both hold. If $S' \neq \varepsilon$, then we immediately have $\Delta(S', R') = \infty$ and $\Delta(S'S'', R'R'') \geq \Delta(S'', R'')$, whence (i) and (ii), respectively.

(iii) We have

$$\begin{aligned}
\mathrm{SD}(S'S'', R'R'') &= \max_i \Delta((S'S'')_{\geq i}, (R'R'')_{\geq i}) \\
&= \max\{\max_i \Delta(S''_{\geq i}, R''_{\geq i}), \max_i \Delta(S'_{\geq i}S'', R'_{\geq i}R'')\} \\
&\leq \max\{\max_i \Delta(S''_{\geq i}, R''_{\geq i}), \max_i \Delta(S'_{\geq i}, R'_{\geq i}), \Delta(S'', R'')\} \\
&= \max\{\max_i \Delta(S''_{\geq i}, R''_{\geq i}), \max_i \Delta(S'_{\geq i}, R'_{\geq i})\} \\
&= \max\{\mathrm{SD}(S'', R''), \mathrm{SD}(S', R')\},
\end{aligned}$$

where the third step uses (i).

(iv) The proof is similar to the previous item:

$$\begin{aligned}
\mathrm{SD}_k(S'S'', R'R'') &= \max_i\{\Delta((S'S'')_{\geq i}, (R'R'')_{\geq i}) : \overline{*}((S'S'')_{<i}) < k\} \\
&= \max_i\{\Delta(S'_{\geq i}S'', R'_{\geq i}R'') : \overline{*}(S'_{<i}) < k\} \\
&\leq \max_i\{\max\{\Delta(S'_{\geq i}, R'_{\geq i}), \Delta(S'', R'')\} : \overline{*}(S'_{<i}) < k\} \\
&= \max\{\mathrm{SD}_k(S', R'), \Delta(S'', R'')\},
\end{aligned}$$

where the second step is valid because $k \leq \overline{*}(S')$ and in particular $i \leq |S'|$, whereas the third step uses (i). □

The following generic lower bound on suffix distance will also be useful.

Proposition 4.13. *Let $k > 0$ be given. Then for all $r \in \Sigma^*$ and $s \in \Sigma^+$,*

$$\mathrm{SD}_k(s, r) \geq 1 - \frac{|r|}{|s|}. \tag{4.3.1}$$

Proof. Fix an arbitrary alignment $S \parallel R$ for s, r. Then

$$\begin{aligned}
\mathrm{SD}_k(S, R) &\geq \Delta(S, R)\\
&= \frac{*(S) + *(R)}{\overline{*}(S)}\\
&= \frac{*(S) + *(R)}{|s|}\\
&= \frac{*(S) + *(S) + |s| - |r|}{|s|}\\
&\geq \frac{|s| - |r|}{|s|},
\end{aligned}$$

where the next-to-last step uses $*(S) + |s| = *(R) + |r|$. □

4.3.2. Existence of Good Tree Codes. The purpose of this section is to prove Theorem 4.9 on the existence of α-good codes, which we now restate for the reader's convenience.

THEOREM 4.14. *For any alphabet* Σ_{in}, *any* $0 < \alpha < 1$, *and any integer* $n \geq 0$, *there is an* α*-good code* $C \colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ *of depth* n *with*

$$|\Sigma_{\text{out}}| = \left\lceil \frac{(10|\Sigma_{\text{in}}|)^{1/\alpha}\, e}{\alpha} \right\rceil^2. \tag{4.3.2}$$

Our treatment is a reworked and simplified version of an argument of Braverman et al. [**25**], who proved the existence of a closely related family of tree codes.

We fix α for the rest of the proof and define Σ_{out} to be the alphabet of consecutive natural numbers, with cardinality given by (4.3.2). For strings u and v, we write $u \diamond v$ to mean that $\mathrm{ED}(u, v) < (1 - \alpha)(|u| + |v|)$. For a tree code $C \colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ of depth n and a string $u \in \Sigma_{\text{in}}^*$, we let C_u denote the tree code $C_u \colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ of depth $n - |u|$ given by $C_u(v) = (C(uv))_{>|u|}$.

Our proof centers around two inductively defined families $\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_n, \ldots$ and $\mathcal{C}_0^*, \mathcal{C}_1^*, \mathcal{C}_2^*, \ldots, \mathcal{C}_n^*, \ldots$, where $\mathcal{C}_n$ and $\mathcal{C}_n^*$ are sets of tree codes of depth n. As a base case, we let $\mathcal{C}_0 = \mathcal{C}_0^*$ be the family whose only member is the tree code $\varepsilon \mapsto \varepsilon$, which is by definition the only tree code of depth 0. Assuming that $\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \ldots, \mathcal{C}_{n-1}$ and $\mathcal{C}_0^*, \mathcal{C}_1^*, \mathcal{C}_2^*, \ldots, \mathcal{C}_{n-1}^*$ have been constructed, we define $\mathcal{C}_n$ to be the family of all tree codes $C\colon \Sigma_{\mathrm{in}}^* \to \Sigma_{\mathrm{out}}^*$ of depth n such that $C_\sigma \in \mathcal{C}_{n-1}^*$ for all $\sigma \in \Sigma_{\mathrm{in}}$, and define $\mathcal{C}_n^*$ to be the family of all α-good codes in $\mathcal{C}_n$. To settle Theorem 4.9, it remains to prove that each $\mathcal{C}_n^*$ is nonempty. We will in fact prove the following stronger claim:

$$\frac{|\mathcal{C}_n^*|}{|\mathcal{C}_n|} \geq \frac{1}{2}, \qquad\qquad n = 0, 1, 2, 3, \ldots. \tag{4.3.3}$$

We will argue by induction on n. The base case $n = 0$ is trivial. For the inductive step, fix $n \geq 1$ arbitrarily and assume that $|\mathcal{C}_i^*|/|\mathcal{C}_i| \geq 1/2$ for $i = 0, 1, 2, \ldots, n-1$. A key technical element of our analysis is the following observation.

Claim 4.15. *Let $u, v, w \in \Sigma_{\mathrm{in}}^*$ be given, where $v_{\leq 1} \neq w_{\leq 1}$. Then*

$$\mathop{\mathbf{P}}_{C \in \mathcal{C}_n} [C(uv) \diamond C_u(w)] \leq \left(\frac{1}{5|\Sigma_{\mathrm{in}}|}\right)^{|u|+|v|+|w|}.$$

The hypothesis $v_{\leq 1} \neq w_{\leq 1}$ above amounts to saying that the longest common prefix of v and w is the empty string.

Proof of Claim 4.15. The claim follows from the following derivation, whose steps we will justify shortly:

$$
\begin{aligned}
\mathop{\mathbf{P}}_{C\in\mathcal{C}_n}[C(uv)\diamond C_u(w)] &\leq 2^{|u|}\mathop{\mathbf{P}}_{\substack{z\in\Sigma_{\mathrm{out}}^{|u|}\\ C\in\mathcal{C}_{n-|u|}}}[zC(v)\diamond C(w)] && (4.3.4)\\
&= 2^{|u|}\mathop{\mathbf{P}}_{\substack{z\in\Sigma_{\mathrm{out}}^{|u|}\\ C',C''\in\mathcal{C}_{n-|u|}}}[zC'(v)\diamond C''(w)] && (4.3.5)\\
&\leq 2^{|u|+|v|}\mathop{\mathbf{P}}_{\substack{z\in\Sigma_{\mathrm{out}}^{|u|}\\ z'\in\Sigma_{\mathrm{out}}^{|v|}\\ C''\in\mathcal{C}_{n-|u|}}}[zz'\diamond C''(w)] && (4.3.6)\\
&\leq 2^{|u|+|v|+|w|}\mathop{\mathbf{P}}_{\substack{z\in\Sigma_{\mathrm{out}}^{|u|}\\ z'\in\Sigma_{\mathrm{out}}^{|v|}\\ z''\in\Sigma_{\mathrm{out}}^{|w|}}}[zz'\diamond z''] && (4.3.7)\\
&\leq 2^{|u|+|v|+|w|}\left(\frac{e}{\alpha\sqrt{|\Sigma_{\mathrm{out}}|}}\right)^{\alpha(|u|+|v|+|w|)} && (4.3.8)\\
&\leq \frac{1}{(5|\Sigma_{\mathrm{in}}|)^{|u|+|v|+|w|}}. && (4.3.9)
\end{aligned}
$$

Inequality (4.3.4) is trivially true for $u=\varepsilon$. To verify validity for $|u|\geq 1$, observe that

$$
\begin{aligned}
\mathop{\mathbf{P}}_{C\in\mathcal{C}_n}[C(uv)\diamond C_u(w)] &= \mathop{\mathbf{P}}_{\substack{z_1\in\Sigma_{\mathrm{out}}\\ C\in\mathcal{C}^*_{n-1}}}[z_1C(u_{\geq 2}v)\diamond C_{u_{\geq 2}}(w)]\\
&\leq \mathop{\mathbf{P}}_{\substack{z_1\in\Sigma_{\mathrm{out}}\\ C\in\mathcal{C}_{n-1}}}[z_1C(u_{\geq 2}v)\diamond C_{u_{\geq 2}}(w)]\cdot\frac{|\mathcal{C}_{n-1}|}{|\mathcal{C}^*_{n-1}|}\\
&\leq \mathop{\mathbf{P}}_{\substack{z_1\in\Sigma_{\mathrm{out}}\\ C\in\mathcal{C}_{n-1}}}[z_1C(u_{\geq 2}v)\diamond C_{u_{\geq 2}}(w)]\cdot 2,
\end{aligned}
$$

where the last two steps use use $\mathcal{C}^*_{n-1}\subseteq\mathcal{C}_{n-1}$ and $|\mathcal{C}^*_{n-1}|\geq|\mathcal{C}_{n-1}|/2$. Applying this maneuver an additional $|u|-1$ times settles (4.3.4). The next step, (4.3.5), is valid because the longest common prefix of v and w is the empty string and therefore $C(v)$ and $C(w)$ are independent. Steps (4.3.6) and (4.3.7) can be verified in a manner identical to (4.3.4). The final steps (4.3.8) and (4.3.9) follow from Proposition 4.4 and (4.3.2), respectively. $\square$

Armed with Claim 4.15, we are now in a position to complete the inductive step. Our objective is to show that $|\mathcal{C}_n^*|/|\mathcal{C}_n| \geq 1/2$, or equivalently that a uniformly random code $C \in \mathcal{C}_n$ has an α-violation with probability at most $1/2$. Recall that an α-violation in C is a quadruple of vertices (A, B, D, E) in the tree representation of C with the following properties:

(i) B is the deepest common predecessor of D and E;

(ii) A is a predecessor of B;

(iii) $AD \diamond BE$, where $AD \in \Sigma_{\text{out}}^*$ and $BE \in \Sigma_{\text{out}}^*$ denote the concatenation of the code symbols along the path from A to D and the path from B to E, respectively.

We further deduce that

(i) A is the root;

(ii) $B \neq E$.

The former holds because the codes in $\mathcal{C}_{n-1}^*$ have no α-violations, and the latter follows from Remark 4.7. These structural constraints allow us to identify an α-violation (A, B, D, E) in C in a one-to-one manner with a triple of strings $u, v, w \in \Sigma_{\text{in}}^*$ such that $v_{\leq 1} \neq w_{\leq 1}$, $w \neq \varepsilon$, and $C(uv) \diamond C_u(w)$. Applying the union bound over all such triples u, v, w,

$$\begin{aligned}
&\mathop{\mathbf{P}}_{C\in\mathcal{C}_n}[C \text{ has an } \alpha\text{-violation}]\\
&\qquad\leq \sum_{\substack{u\in\Sigma_{\text{in}}^*:\\ |u|<n}} \sum_{\substack{v\in\Sigma_{\text{in}}^*:\\ |v|\leq n-|u|}} \sum_{\substack{w\in\Sigma_{\text{in}}^+:\\ |w|\leq n-|u|,\\ w_{\leq 1}\neq v_{\leq 1}}} \mathop{\mathbf{P}}_{C\in\mathcal{C}_n}[C(uv) \diamond C_u(w)].
\end{aligned}$$

Appealing to Claim 4.15 and simplifying,

$$\begin{aligned}
&\mathop{\mathbf{P}}_{C\in\mathcal{C}_n}[C \text{ has an } \alpha\text{-violation}]\\
&\qquad\leq \sum_{\substack{u\in\Sigma_{\text{in}}^*:\\ |u|<n}} \sum_{\substack{v\in\Sigma_{\text{in}}^*:\\ |v|\leq n-|u|}} \sum_{\substack{w\in\Sigma_{\text{in}}^+:\\ |w|\leq n-|u|,\\ w_{\leq 1}\neq v_{\leq 1}}} \left(\frac{1}{5|\Sigma_{\text{in}}|}\right)^{|u|+|v|+|w|}\\
&\qquad\leq \sum_{u\in\Sigma_{\text{in}}^*} \sum_{v\in\Sigma_{\text{in}}^*} \sum_{w\in\Sigma_{\text{in}}^+} \left(\frac{1}{5|\Sigma_{\text{in}}|}\right)^{|u|+|v|+|w|}\\
&\qquad= \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \sum_{k=1}^{\infty} \frac{1}{5^{i+j+k}}\\
&\qquad= \frac{1}{5}\cdot\frac{1}{\left(1-\frac{1}{5}\right)^3}\\
&\qquad< \frac{1}{2}.
\end{aligned}$$

The final inequality is equivalent to $|\mathcal{C}_n^*|/|\mathcal{C}_n| > 1/2$, completing the inductive step. We have settled (4.3.3) and thereby proved Theorem 4.9.

4.3.3. Longest prefix decoding. In interactive coding, a sequence of symbols is encoded with a tree code and transmitted over an unreliable channel. On the receiving end, an attempt is then made to decode the sequence. The encoding and decoding are fundamentally different in that the former is fully determined by the tree code, whereas the latter allows for several reasonable approaches. In contrast to the work of Braverman et al. [**25**], our interactive coding schemes use *longest prefix decoding*, whereby the receiver attempts to correctly decode as long a prefix of the original sequence as possible. The following key theorem relates the length of such a prefix to the suffix distance between the original sequence and its received counterpart.

THEOREM 4.16. *Fix an α-good code $C\colon \Sigma_{\mathrm{in}}^* \to \Sigma_{\mathrm{out}}^*$ with $0 < \alpha < 1$. Consider a string $r \in \Sigma_{\mathrm{out}}^*$ and codewords s', s'' of C with*

$$\mathrm{SD}_k(s', r) < 1 - \alpha,$$

$$\mathrm{SD}_k(s'', r) < 1 - \alpha.$$

Then

$$s'_{\leq k} = s''_{\leq k}. \tag{4.3.10}$$

Previous work [**25**] settled a special case of Theorem 4.16 for $k = \infty$, corresponding to the correct decoding of the entire sequence. The extension to arbitrary k is essential to the optimal interactive coding schemes in our work.

Proof of Theorem 4.16 (cf. Braverman et al. [**25**]). Let s be the longest common prefix of s' and s''. For the sake of contradiction, assume that (4.3.10) fails. Then

$$s' \neq s'', \tag{4.3.11}$$

$$|s| < k. \tag{4.3.12}$$

We will show that these two conditions force an α-violation in C, contrary to the theorem hypothesis.

Fix alignments $S' \parallel R'$ and $S'' \parallel R''$ for the string pairs s', r and s'', r, respectively, such that

$$\mathrm{SD}_k(S', R') < 1 - \alpha, \tag{4.3.13}$$

$$\mathrm{SD}_k(S'', R'') < 1 - \alpha. \tag{4.3.14}$$

Let $i', i'' \geq 0$ be integers with

$$s = \not\ast(S'_{\leq i'}), \tag{4.3.15}$$

$$s = \not\ast(S''_{\leq i''}). \tag{4.3.16}$$

It follows from (4.3.11) that

$$\overline{\ast}(S'_{>i'}) + \overline{\ast}(S''_{>i''}) > 0. \tag{4.3.17}$$

Observe also that r contains both $\not\ast(R'_{>i'})$ and $\not\ast(R''_{>i''})$ as suffixes, which means that one of those strings is a suffix of the other. Without loss of generality, assume that $\not\ast(R''_{>i''})$ is a suffix of $\not\ast(R'_{>i'})$ and fix an integer $j'' \geq 0$ such that

$$j'' \leq i'', \tag{4.3.18}$$

$$\not\ast(R''_{>j''}) = \not\ast(R'_{>i'}). \tag{4.3.19}$$

It follows from (4.3.12) and (4.3.15) that $\overline{\ast}(S'_{\leq i'}) < k$. Analogously, (4.3.12), (4.3.16), and (4.3.18) give $\overline{\ast}(S''_{\leq j''}) < k$. Therefore, the suffix distance bounds (4.3.13) and (4.3.14) guarantee that

$$\Delta(S'_{>i'}, R'_{>i'}) < 1 - \alpha, \tag{4.3.20}$$

$$\Delta(S''_{>j''}, R''_{>j''}) < 1 - \alpha. \tag{4.3.21}$$

In addition, (4.3.17) and (4.3.18) imply that

$$\overline{*}(S'_{>i'}) + \overline{*}(S''_{>j''}) > 0. \tag{4.3.22}$$

Now

$$\begin{aligned}
&\mathrm{ED}(\not{*}(S'_{>i'}), \not{*}(S''_{>j''})) \\
&\qquad \leq \mathrm{ED}(\not{*}(S'_{>i'}), \not{*}(R'_{>i'})) + \mathrm{ED}(\not{*}(R'_{>i'}), \not{*}(S''_{>j''})) \\
&\qquad = \mathrm{ED}(\not{*}(S'_{>i'}), \not{*}(R'_{>i'})) + \mathrm{ED}(\not{*}(S''_{>j''}), \not{*}(R'_{>i'})) \\
&\qquad = \mathrm{ED}(\not{*}(S'_{>i'}), \not{*}(R'_{>i'})) + \mathrm{ED}(\not{*}(S''_{>j''}), \not{*}(R''_{>j''})) \\
&\qquad \leq \Delta(S'_{>i'}, R'_{>i'}) \cdot \overline{*}(S'_{>i'}) + \Delta(S''_{>j''}, R''_{>j''}) \cdot \overline{*}(S''_{>j''}) \\
&\qquad < (1-\alpha)(\overline{*}(S'_{>i'}) + \overline{*}(S''_{>j''})),
\end{aligned} \tag{4.3.23}$$

where the first four steps follow from (4.2.3), (4.2.2), (4.3.19), and Fact 4.5, respectively, and the final step is immediate from (4.3.20)–(4.3.22).

It remains to interpret our findings in terms of the tree representation of C. Let A, B, D, E be the vertices reached by following the paths $\not{*}(S''_{\leq j''}), s, s'', s'$, respectively, from the root of the tree. Then (4.3.23) is equivalent to $\mathrm{ED}(BE, AD) < (1-\alpha)(|BE| + |AD|)$, which is the promised α-violation. □

We are now in a position to describe our decoding algorithm and relate its decoding guarantees to the suffix distance between the original sequence and its received counterpart.

THEOREM 4.17. *Let $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ be an α-good code, $0 < \alpha < 1$. Then there is an algorithm $\textsc{Decode}_{C,\alpha}\colon \Sigma_{\text{out}}^* \to \Sigma_{\text{out}}^*$ that runs in finite time and obeys*

$$(\textsc{Decode}_{C,\alpha}(r))_{\leq k} = s_{\leq k} \tag{4.3.24}$$

for any real $0 \le k \le \infty$, *any codeword* s, *and any string* $r \in \Sigma_{\mathrm{out}}^*$ *with* $\mathrm{SD}_k(s,r) < 1-\alpha$.

Proof. For a codeword s and a string r, define

$$K(s,r) = \max\{k \in \mathbb{N} \cup \{\infty\} : \mathrm{SD}_k(s,r) < 1-\alpha\}.$$

The maximization on the right-hand side is over a nonempty set that contains $k=0$, so that $K(s,r)$ is well-defined for every s, r pair. The algorithm is the natural one: on input r, the output of $\textsc{Decode}_{C,\alpha}$ is any $s^* \in \arg\max_s K(s,r)$, where s ranges over all codewords of C. To verify (4.3.24), let s be any codeword with $\mathrm{SD}_k(s,r) < 1-\alpha$. Then the algorithm output s^* obeys $\mathrm{SD}_k(s^*,r) < 1-\alpha$ and hence $s^*_{\le k} = s_{\le k}$ by Theorem 4.16.

It remains to show that $\textsc{Decode}_{C,\alpha}$ can be implemented to run in finite time. Clearly, computing $K(s,r)$ for any pair of strings s and r takes finite time. To find a codeword in $\arg\max_s K(s,r)$, it is suffices to consider codewords of length at most r/α because longer codewords s satisfy $K(s,r) = 0$ by Proposition 4.13. □

4.3.4. Frequency of good decodings. In the analysis of interactive coding schemes, one typically needs to argue that there are *many* points in time when the receiving party is able to correctly decode the sequence of symbols transmitted so far. We estimate the number of such "good decodings" using the following technical fact, closely analogous to previous work [**26, 25**].

Proposition 4.18. *Fix an alignment* $S \parallel R$ *and define*

$$G = \{i : S_i = R_i \ne *\},$$

$$D = \{i : S_i \ne *, R_i = *\},$$

$$I = \{i : S_i = *, R_i \ne *\}.$$

```
$A \leftarrow \varnothing$
$i \leftarrow \ell$
while $i > 0$ do
    if $\mathrm{SD}(S_1 S_2 \dots S_i, R_1 R_2 \dots R_i) < 1 - \alpha$ then
        $A \leftarrow A \cup \{i\}$
        $i \leftarrow i - 1$
    else
        find any index $j$ with $\Delta(S_j S_{j+1} \dots S_i, R_j R_{j+1} \dots R_i) \geq 1 - \alpha$
        $i \leftarrow j - 1$
    end
end
return $A$
```

Algorithm 1: An algorithm to accompany the proof of Proposition 4.18.

Then for all $0 < \alpha < 1$,

$$|\{i \in G : \mathrm{SD}(S_1 S_2 \dots S_i, R_1 R_2 \dots R_i) < 1 - \alpha\}| \\ \geq |G| - \frac{\alpha}{1-\alpha}|D| - \frac{1}{1-\alpha}|I|.$$

The notation in Proposition 4.18 is mnemonic, with I, D, and G denoting the positions of the inserted, deleted, and "good" (unchanged) symbols, respectively. Note that insertions and deletions play asymmetric roles in this result, insertions being more damaging than deletions.

Proof of Proposition 4.18 (adapted from [**26, 25**]). Abbreviate $\ell = |S| = |R|$ and consider Algorithm 1, which iteratively constructs a subset

$$A \subseteq \{i : \mathrm{SD}(S_1 S_2 \dots S_i, R_1 R_2 \dots R_i) < 1 - \alpha\}. \tag{4.3.25}$$

Since $\mathrm{SD}(S_1 S_2 \dots S_i, R_1 R_2 \dots R_i) \geq \Delta(S_i, R_i) \geq 1$ for every $i \in I \cup D$, we infer that

$A \subseteq G$. In particular,

$$\Delta(S_{\overline{A}}, R_{\overline{A}}) = \frac{|I \cap \overline{A}| + |D \cap \overline{A}|}{|G \cap \overline{A}| + |D \cap \overline{A}|}$$
$$= \frac{|I| + |D|}{|G| - |A| + |D|}. \tag{4.3.26}$$

The complementary set $\overline{A}$ is the disjoint union of the intervals $\{j, j+1, \ldots, i\}$ computed by the **else** clause, each of which satisfies $\Delta(S_j S_{j+1} \ldots S_i, R_j R_{j+1} \ldots R_i) \geq 1 - \alpha$. It follows by Proposition 4.12(ii) that $\Delta(S_{\overline{A}}, R_{\overline{A}}) \geq 1 - \alpha$, which along with (4.3.26) gives

$$|A| \geq |G| - \frac{\alpha}{1-\alpha}|D| - \frac{1}{1-\alpha}|I|.$$

In view of (4.3.25) and $A \subseteq G$, the proof is complete. □

4.4. A coding scheme with a polynomial-size alphabet

We will now show how to faithfully simulate any protocol in the adversarial setting at the expense of a large increase in alphabet size and a constant-factor increase in communication cost. For an arbitrary constant $\varepsilon > 0$, we give an interactive coding scheme that tolerates corruption rate $\frac{1}{4} - \varepsilon$ as well as any normalized corruption rate $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$ with $\varepsilon_{\text{subs}} + \frac{3}{4}\varepsilon_{\text{oos}} \leq \frac{1}{4} - \varepsilon$. In detail, the main result of this section is as follows.

THEOREM 4.19. *Fix an arbitrary constant $\varepsilon > 0$, and let π be an arbitrary protocol with alphabet Σ. Then there exists an interactive coding scheme for π with alphabet size $(|\Sigma| \cdot |\pi|)^{O(1)}$ and communication cost $O(|\pi|)$ that tolerates*

(i) *corruption rate $\frac{1}{4} - \varepsilon$;*

(ii) *any normalized corruption rate $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$ with $\varepsilon_{\text{subs}} + \frac{3}{4}\varepsilon_{\text{oos}} \leq \frac{1}{4} - \varepsilon$.*

As we will see later in this chapter, Theorem 4.19 is optimal with respect to the corruption rate and normalized corruption rate that it tolerates. We have organized our proof of the theorem around nine milestones, corresponding to Sections 4.4.1–4.4.9. Looking ahead, we will obtain the main result of this chapter by improving the alphabet size to a constant.

4.4.1. The simulation. Recall that any protocol can be brought into canonical form at the expense of doubling its communication cost. We may therefore assume that π is in canonical form to start with. As a result, we may identify Alice's input with a set X of odd-depth edges of the protocol tree for π, and Bob's input with a set Y of even-depth edges. Execution of π corresponds to a walk down the unique root-to-leaf path in $X \cup Y$, whose length we denote by

$$n = |\pi|.$$

Analogous to previous work [**26, 25**], our interactive coding scheme involves Alice and Bob sending edges from their respective input sets X and Y. At any given point, Alice will send an edge e only if she has already sent every proper predecessor of e in X, and likewise for Bob. This makes it possible for the sender to represent an edge e succinctly as a pair (i, σ), where i is the index of a previous transmission by the sender that featured the grandparent of e, and $\sigma \in \Sigma \times \Sigma$ uniquely identifies e relative to that grandparent. When transmitting an edge e of depth 1 or 2, the sender sets $i = 0$ to indicate that there are no proper predecessors to refer to. Viewing each (i, σ) pair as an alphabet symbol, the resulting alphabet Σ_{in} has size at most $|\Sigma|^2$ multiplied by the total number of transmissions. The following lemma shows that given any sequence of edge representations, it is always possible to recover the corresponding sequence of edges.

```
Input: $X$ (set of Alice's edges)
encode and send the edge in $X$ incident to the root
foreach $i = 1, 2, 3, \ldots, N$ do
    receive a symbol $r_i \in \Sigma_{\text{out}}$
    $s \leftarrow \textsc{Decode}_{C,\alpha}(r_1 r_2 \ldots r_i)$
    interpret $s$ as a sequence $B$ of even-depth edges
    $\ell \leftarrow$ maximum length of a rooted path in $X \cup B$
    compute the shortest prefix of $B$ s.t. $X \cup B$ contains a rooted path of length $\ell$,
      and let $P$ be the rooted path so obtained
    $out \leftarrow$ deepest vertex in $P$
    if $i \leq N - 1$ then
        encode and send the deepest edge in $P \cap X$ whose proper predecessors in $X$
          have all been sent
    end
end
```

Algorithm 2: Coding scheme for Alice

Lemma 4.20. *Consider an arbitrary point in time, and let*

$$(i_1, \sigma_1), (i_2, \sigma_2), \ldots (i_t, \sigma_t) \tag{4.4.1}$$

be the sequence of edge representations sent so far by one of the players. Then the sequence uniquely identifies the corresponding edges $e_1, e_2, \ldots, e_t$ sent by that player.

Proof. The proof is by induction of t, the base case $t = 0$ being trivial. For the inductive step, let $e_1, e_2, \ldots, e_{t-1}$ be the unique sequence of edges corresponding to $(i_1, \sigma_1), (i_2, \sigma_2), \ldots, (i_{t-1}, \sigma_{t-1})$. Recall that $i_t \in \{0, 1, 2, \ldots, t-1\}$. If $i_t \in \{1, 2, \ldots, t-1\}$, then by definition (i_t, σ_t) is the grandchild of e_{i_t} that corresponds to $\sigma_t \in \Sigma \times \Sigma$. If $i_t = 0$, then by definition (i_t, σ_t) is the edge of depth 1 in Alice's case, or depth 2 in Bob's, that corresponds to σ_t. □

```
Input: $Y$ (set of Bob's edges)
1 foreach $i = 1, 2, 3, \ldots, N$ do
2     receive a symbol $r_i \in \Sigma_{\text{out}}$
3     $s \leftarrow \textsc{Decode}_{C,\alpha}(r_1 r_2 \ldots r_i)$
4     interpret $s$ as a sequence $A$ of odd-depth edges
5     $\ell \leftarrow$ maximum length of a rooted path in $Y \cup A$
6     compute the shortest prefix of $A$ s.t. $Y \cup A$ contains a rooted path of length $\ell$,
        and let $P$ be the rooted path so obtained
7     $out \leftarrow$ deepest vertex in $P$
8     encode and send the deepest edge in $P \cap Y$ whose proper predecessors in $Y$
        have all been sent
9 end
```

Algorithm 3: Coding scheme for Bob

A formal description of the interactive coding scheme is given by Algorithms 2 and 3 for Alice and Bob, respectively. In this description, $\alpha = \alpha(\varepsilon) \in (0, 1)$ and $N = N(n, \varepsilon)$ are parameters to be set later, and $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ is an arbitrary α-good code whose existence is ensured by Theorem 4.9. Alice and Bob use C to encode every transmission. In particular, the encoded symbol from Σ_{out} at any given point depends not only on the symbol from Σ_{in} being transmitted but also on the content of previous transmissions by the sender. The decoding is done using the $\textsc{Decode}_{C,\alpha}$ algorithm of Theorem 4.17. Apart from the initial send by Alice in line 1, the roles of two players are symmetric. In particular, the pseudocode makes it clear that Alice and Bob send at most N transmissions each. We conclude that $|\Sigma_{\text{in}}| \leq |\Sigma|^2 \cdot 2N$ and therefore by Theorem 4.9,

$$|\Sigma_{\text{out}}| = (|\Sigma| \cdot N)^{O(1/\alpha)}. \tag{4.4.2}$$

We pause to elaborate on the decoding and interpretation steps in lines 4–5 for Alice and lines 3–4 for Bob. The decoding step produces a codeword s of C, which by Fact 4.8 corresponds to a unique string in Σ_{in}^*. Recall that this string is of the form (4.4.1) for some integers $i_1, i_2, \ldots, i_t$ and some $\sigma_1, \sigma_2, \ldots, \sigma_t \in \Sigma \times \Sigma$. The receiving party uses the inductive procedure of Lemma 4.20 to convert (4.4.1) to a sequence of edges. It may happen that (4.4.1) is syntactically malformed; in that case, the receiving party interrupts the interpretation process at the longest prefix of (4.4.1) that corresponds to a legitimate sequence of edges. This completes the interpretation step, yielding a sequence of edges A for Bob and B for Alice.

In Sections 4.4.2–4.4.9 below, we examine an arbitrary but fixed execution of the interactive coding scheme. In particular, we will henceforth consider the inputs X and Y and the adversary's actions to be fixed. We allow any behavior by the adversary as long as it meets one of the criteria (i), (ii) in Theorem 4.19. We will show that as soon as the communication stops, the variable *out* is set for both Alice and Bob to the leaf vertex of the unique root-to-leaf path in $X \cup Y$. This will prove Theorem 4.19.

4.4.2. Events. A central notion in our analysis is that of an *event.* There are three types of events: deletions, insertions, and good events. A successful transmission corresponds to a single event, which we call a *good event.* A transmission that is subject to an attack, on the other hand, corresponds to two events, namely, a *deletion event* followed immediately by an *insertion event.* Each event has an *addressee.* The addressee of a good event is defined to be the receiver of the transmission. Similarly, the deletion and insertion events that arise from a substitution attack are said to be addressed to the receiver of the transmission. In an out-of-sync attack, on the other hand, the deletion event is addressed to the intended receiver of the transmission, whereas the insertion event is addressed to the sender.

To illustrate these definitions, consider the hypothetical execution in Table 1. The columns of the table are numbered 1 through 10, corresponding the ten transmissions

Transmission #	**1**	**2**		**3**	**4**	**5**	**6**		**7**	**8**		**9**	**10**
Addressee	B	A	A	B	A	B	A	B	A	B	A	B	A
Symbol sent	0	1	$*$	0	1	0	0	$*$	1	0	$*$	0	1
Symbol received	0	$*$	0	0	1	0	$*$	0	1	$*$	1	0	1

Table 1. A hypothetical execution.

sent in this execution. These ten columns are further split into subcolumns that correspond to individual events, as follows.

(i) Transmissions $1, 3, 4, 5, 7, 9, 10$ result in successful deliveries, each contributing a good event addressed to the receiver of the transmission. For each of these transmissions, the entries in the sent and received rows coincide and show the symbol delivered from the sender to the receiver.

(ii) Transmission 2 is subject to a substitution attack, whereby the sent symbol "1" is deleted (corresponding to the "1" and $*$ entries in the sent and received rows, respectively) and a symbol of "0" is inserted in its place (corresponding to the $*$ and "0" entries in the sent and received rows, respectively). Transmission 2 thus contributes a deletion event and an insertion event, both addressed to the receiver of the transmission.

(iii) Transmissions 6 and 8 are subject to out-of-sync attacks, each contributing a deletion event and an insertion event. In both cases, the deletion event is addressed to the transmission's intended receiver, whereas the insertion event is addressed to the transmission's sender. In the case of transmission 6, the sent symbol "0" is deleted (corresponding to the "0" and $*$ entries in the sent and received rows, respectively) and a new symbol of "0" is spuriously sent back on behalf of the transmission's intended receiver (corresponding to the $*$ and "0" entries in the sent and received rows, respectively).

Execution of the interactive coding scheme gives rise to a string alignment $S' \parallel R'$ for Alice and a string alignment $S'' \parallel R''$ for Bob. Each position i in the strings S' and R'

corresponds in a one-to-one manner to an event addressed to Alice, which is either a good event ($S'_i = R'_i$), a deletion ($S'_i \neq *, R'_i = *$), or an insertion ($S'_i = *, R'_i \neq *$). An analogous description applies to Bob's strings S'' and R''. To illustrate, the execution in Table 1 corresponds to

$$\begin{aligned} S' &= 1{*}101{*}1, \\ R' &= {*}01{*}111 \end{aligned}$$

and

$$\begin{aligned} S'' &= 000{*}00, \\ R'' &= 0000{*}0. \end{aligned}$$

For integers $i \leq j$, we let $S'[i,j] \parallel R'[i,j]$ denote the subalignment of $S' \parallel R'$ that corresponds to transmissions $i, i+1, \ldots, j$. Analogously, $S''[i,j] \parallel R''[i,j]$ denotes the subalignment of $S'' \parallel R''$ that corresponds to transmissions $i, i+1, \ldots, j$. We alert the reader that in our notation, S'_i and $S'[i,i]$ have completely different meanings: the former is the i^{th} symbol of S', whereas the latter is the substring of S' that corresponds to the i^{th} transmission. We define

$$\begin{aligned} G' &= \{i : S'[i,i] = R'[i,i] \neq \varepsilon\}, \\ D' &= \{i : R'[i,i] = *\}, \\ I' &= \{i : S'[i,i] = *\}. \end{aligned}$$

In words, G', D', I' are the sets of transmissions that contribute a good event, a deletion event, and an insertion event, respectively, addressed in each case to Alice.

We define analogous sets for Bob:

$$G'' = \{i : S''[i,i] = R''[i,i] \neq \varepsilon\},$$

$$D'' = \{i : R''[i,i] = *\},$$

$$I'' = \{i : S''[i,i] = *\}.$$

We abbreviate

$$G = G' \cup G'',$$

$$D = D' \cup D'',$$

$$I = I' \cup I''.$$

We let T denote the combined number of transmissions sent by Alice and Bob. Since neither player can send more than N transmissions, we have

$$T \leq 2N. \tag{4.4.3}$$

The following lemma collects basic properties of the sets just introduced.

Lemma 4.21. *The following properties hold:*

(i) G' *and* G'' *form a partition of* G;
(ii) I' *and* I'' *form a partition of* I;
(iii) D' *and* D'' *form a partition of* D;
(iv) $I = D$;
(v) $I' \setminus D' = D'' \setminus I''$;
(vi) $I'' \setminus D'' = D' \setminus I'$;
(vii) G *and* I *form a partition of* $\{1, 2, \ldots, T\}$;
(viii) G *and* D *form a partition of* $\{1, 2, \ldots, T\}$.

Proof. Properties (i)–(iii) hold because any given transmission contributes at most one good event, at most one deletion event, and at most one insertion event, where each event is addressed to precisely one of the players. Property (iv) holds because deletions and insertions always occur in pairs, with any given transmission generating both or neither. Property (v) follows set-theoretically from the preceding properties:

$$\begin{aligned}
I' \setminus D' &= (I \setminus I'') \setminus (D \setminus D'') && \text{by (ii) and (iii)}\\
&= (D \setminus I'') \setminus (D \setminus D'') && \text{by (iv)}\\
&= D \cap \overline{I''} \cap \overline{D \cap \overline{D''}} && \text{by Boolean algebra}\\
&= D \cap \overline{I''} \cap (\overline{D} \cup D'') && \text{by Boolean algebra}\\
&= D \cap \overline{I''} \cap D'' && \text{by Boolean algebra}\\
&= \overline{I''} \cap D'' && \text{by (iii)}\\
&= D'' \setminus I'' && \text{by Boolean algebra.}
\end{aligned}$$

The proof of (vi) is entirely analogous. Properties (vii) and (viii) can be restated by saying that the transmissions can be partitioned into those that result in successful deliveries and those that are subject to an attack. □

4.4.3. Excellent transmissions. A straightforward consequence of Lemma 4.20 is that the codewords $\not\ast(S'[1,t])$ and $\not\ast(S''[1,t])$ completely reveal the sequences of edges sent by Bob and by Alice, respectively, over the course of the first t transmissions. We formalize this observation below.

Lemma 4.22. *Let $t \in \{1, 2, \ldots, T\}$ be given. Then:*

(i) *the string $\not\ast(S'[1,t])$ uniquely identifies the sequence of protocol tree edges that Bob sends Alice over the course of transmissions $1, 2, \ldots, t$;*

(ii) *the string* $\not*(S''[1,t])$ *uniquely identifies the sequence of protocol tree edges that Alice sends Bob over the course of transmissions* $1, 2, \ldots, t$.

Proof. By symmetry, it suffices to prove the former claim. By Fact 4.8, the codeword $\not*(S'[1,t]) \in \Sigma_{\text{out}}^*$ corresponds to a unique string in Σ_{in}^*, which is the sequence of edge representations that Bob sent Alice over the course of the first t transmissions. By Lemma 4.20, this sequence of edge representations uniquely identifies the edges themselves. □

Of course, due to interference by the adversary, the receiving party rarely if ever has access to the exact codeword sent by his or her counterpart. This motivates us to identify sufficient conditions that allow for complete and correct decoding by the receiving party. Define

$$E' = \{i \in G' : \mathrm{SD}(S'[1,i], R'[1,i]) < 1 - \alpha\},$$

$$E'' = \{i \in G'' : \mathrm{SD}(S''[1,i], R''[1,i]) < 1 - \alpha\}.$$

We refer to E' and E'' as the sets of *excellent* transmissions for Alice and Bob, respectively. This term is borne out by the following lemma.

Lemma 4.23. *Let* $t \in \{1, 2, \ldots, T\}$ *be given.*

(i) *If* $t \in E'$, *then on receipt of transmission* t, *Alice is able to correctly recover the complete sequence of edges that Bob has sent her by that time.*

(ii) *If* $t \in E''$, *then on receipt of transmission* t, *Bob is able to correctly recover the complete sequence of edges that Alice has sent him by that time.*

Proof. By symmetry, it again suffices to prove the former claim. Let $t \in E'$. Then by definition, $\mathrm{SD}(S'[1,t], R'[1,t]) < 1 - \alpha$. Taking $k = \infty$ in Theorem 4.17, we conclude that $\textsc{Decode}_{C,\alpha}(\not*(R'[1,t])) = \not*(S'[1,t])$. This means that on receipt of transmission

t, Alice is able to correctly recover the entire codeword $\not*(S'[1,t])$ that Bob has sent her so far. By Lemma 4.22, this in turn makes it possible for Alice to correctly identify the corresponding sequence of edges. $\square$

4.4.4. Bad transmissions. Recall that each symbol received by Alice from the communication channel corresponds in a one-to-one manner to a good event or an insertion. Put another way, each such symbol originates in a one-to-one manner from a transmission in $G' \cup I'$. As we saw in Section 4.4.3, the symbols that correspond to excellent transmissions $E' \subseteq G' \cup I'$ allow Alice to correctly recover the sequence of edges that Bob has sent her so far. In all other cases, the conversion of the received string to an edge sequence can produce unpredictable results and cannot be trusted. This motivates us to define the sets of *bad* transmissions for Alice and Bob by

$$\begin{aligned} B' &= (G' \cup I') \setminus E', \\ B'' &= (G'' \cup I'') \setminus E'', \end{aligned}$$

respectively. We abbreviate

$$B = B' \cup B''.$$

Lemma 4.24. *The sets B' and B'' form a partition of B.*

Proof:

$$\begin{aligned} B' \cap B'' &\subseteq (G' \cup I') \cap (G'' \cup I'') \\ &= (G' \cap G'') \cup (I' \cap I'') \cup (G' \cap I'') \cup (G'' \cap I') \\ &\subseteq (G' \cap G'') \cup (I' \cap I'') \cup (G \cap I) \\ &= \varnothing, \end{aligned}$$

where the last step follows from Lemma 4.21 (i), (ii), (vii). □

As one might expect, the number of bad transmissions is closely related to the number of attacks by the adversary. This relation is formalized by the following lemma.

LEMMA 4.25. *For any interval J with $1 \in J$,*

$$|B|_J \leq \frac{2}{1-\alpha}|D|_J.$$

The reader will recall that $|B|_J = |B \cap J|$ and $|D|_J = |D \cap J|$ in the lemma above. We use this relative cardinality notation extensively in the rest of the chapter for improved readability and ease of typesetting.

Proof of Lemma 4.25. Since B and D are sets of positive integers, it suffices to consider an *integer* interval $J = \{1, 2, \ldots, t\}$. Applying Proposition 4.18 to the alignment $S'[1,t] \parallel R'[1,t]$ shows that

$$\begin{aligned} |E' \cap \{1,2,\ldots,t\}| \geq{} & |G' \cap \{1,2,\ldots,t\}| \\ & - \frac{\alpha}{1-\alpha}|D' \cap \{1,2,\ldots,t\}| - \frac{1}{1-\alpha}|I' \cap \{1,2,\ldots,t\}|, \end{aligned}$$

which can be succinctly written as

$$|E'|_J \geq |G'|_J - \frac{\alpha}{1-\alpha}|D'|_J - \frac{1}{1-\alpha}|I'|_J. \tag{4.4.4}$$

Now

$$\begin{aligned}
|B'|_J &= |(G' \cup I') \setminus E'|_J \\
&= |G' \cup I'|_J - |E'|_J \\
&= |G'|_J + |I'|_J - |E'|_J \\
&\leq \frac{\alpha}{1-\alpha}|D'|_J + \frac{2-\alpha}{1-\alpha}|I'|_J, \qquad (4.4.5)
\end{aligned}$$

where the first step holds by definition, the second uses the containment $E' \subseteq G'$, the third is valid by Lemma 4.21 (vii), and the fourth follows from (4.4.4). A symmetric argument gives

$$|B''|_J \leq \frac{\alpha}{1-\alpha}|D''|_J + \frac{2-\alpha}{1-\alpha}|I''|_J. \qquad (4.4.6)$$

As a result,

$$\begin{aligned}
|B|_J &\leq \frac{\alpha}{1-\alpha}(|D'|_J + |D''|_J) + \frac{2-\alpha}{1-\alpha}(|I'|_J + |I''|_J) \\
&= \frac{\alpha}{1-\alpha}|D|_J + \frac{2-\alpha}{1-\alpha}|I|_J \\
&= \frac{2}{1-\alpha}|D|_J,
\end{aligned}$$

where the first step follows from (4.4.5) and (4.4.6), the second uses Lemma 4.21 (ii), (iii), and the third uses Lemma 4.21 (iv). □

4.4.5. Virtual length. Key to our approach is a virtual view of communication that centers around *events* rather than actual transmissions. In particular, we focus on alternations in event addressee as opposed to alternations in sender. To start with, we define for an arbitrary set $Z \subseteq \mathbb{R}$ its *virtual length* by

$$\tilde{\|}Z\tilde{\|} = |G' \cup I' \cup D'|_Z + |G'' \cup I'' \cup D''|_Z. \qquad (4.4.7)$$

In other words, the virtual length $\tilde{\|}Z\tilde{\|}$ is the number of transmissions in Z that have an event addressed to Alice, plus the number of transmissions in Z that have an event addressed to Bob. It follows immediately that

$$|Z| \leq \tilde{\|}Z\tilde{\|} \leq 2|Z|$$

for any $Z \subseteq \{1, 2, \ldots, T\}$, and a moment's thought reveals that the lower and upper bounds can both be attained. We are of course interested only in subsets $Z \subseteq \{1, 2, \ldots, T\}$, but defining virtual length as we did above for arbitrary $Z \subseteq \mathbb{R}$ greatly simplifies the notation. We now show that in the special case when Z is an interval, the two summands in (4.4.7) differ by at most 1.

Lemma 4.26. *For any interval J,*

$$\tilde{\|}J\tilde{\|} \leq 2|G' \cup I' \cup D'|_J + 1, \tag{4.4.8}$$

$$\tilde{\|}J\tilde{\|} \leq 2|G'' \cup I'' \cup D''|_J + 1 \tag{4.4.9}$$

and

$$\tilde{\|}J\tilde{\|} \geq 2|G' \cup I' \cup D'|_J - 1, \tag{4.4.10}$$

$$\tilde{\|}J\tilde{\|} \geq 2|G'' \cup I'' \cup D''|_J - 1. \tag{4.4.11}$$

Proof. Consider arbitrary integers $i_1 < i_2$ such that

$$i_1 \in (G'' \cup D'' \cup I'') \setminus (G' \cup D' \cup I'),$$

$$i_2 \in (G'' \cup D'' \cup I'') \setminus (G' \cup D' \cup I').$$

The first equation states that transmission i_1 is sent by Alice and is not subject to an out-of-sync attack. Recall that a transmission causes a change of speaker if and only if it is not subject to an out-of-sync attack. As a result, a change of speaker from

Alice to Bob happens immediately after transmission i_1. Since the later transmission i_2 is again sent by Alice, there must be an intermediate transmission j that causes a change of speaker from Bob to Alice. This implies

$$j \in (G' \cup D' \cup I') \setminus (G'' \cup D'' \cup I'').$$

The previous paragraph shows that the interval between any two distinct integers in $(G''\cup D''\cup I'')\setminus(G'\cup D'\cup I')$ contains at least one integer in $(G'\cup D'\cup I')\setminus(G''\cup D''\cup I'')$. We conclude that for any interval J,

$$|G'' \cup D'' \cup I''|_J \leq |G' \cup D' \cup I'|_J + 1.$$

Adding $|G' \cup D' \cup I'|_J$ to both sides of this inequality yields (4.4.8), whereas adding $|G'' \cup D'' \cup I''|_J$ to both sides yields (4.4.11). A symmetric argument settles the remaining inequalities (4.4.9) and (4.4.10). □

We now show that the combined virtual length of all transmissions is at least $2N$. This contrasts with the number of transmissions themselves, which can be significantly less than $2N$ due to out-of-sync attacks.

Lemma 4.27. *The total virtual length of all transmissions satisfies*

$$\widetilde{\|}[1, T]\widetilde{\|} \geq 2N.$$

Proof. For the communication to stop, one of the players needs to terminate. This happens only when that player has sent N symbols and received as many. Formulaically, this translates to

$$\begin{aligned} |G'' \cup D''| &\geq N, \\ |G' \cup I'| &\geq N \end{aligned}$$

if Alice terminates first, and

$$|G' \cup D'| \geq N,$$
$$|G'' \cup I''| \geq N$$

if Bob terminates first. Either way,

$$\begin{aligned} \tilde{\|}[1,T]\tilde{\|} &= |G' \cup D' \cup I'| + |G'' \cup D'' \cup I''| \\ &\geq 2N. \end{aligned}$$ □

Next, we relate the virtual length of any interval to the number of attacks experienced by Alice and Bob during that time.

LEMMA 4.28. *Let i, j be given integers with $i \leq j$. Then*

$$\tilde{\|}[i,j]\tilde{\|} \leq \frac{4|D|_{[i,j]}}{\delta} + 1 \tag{4.4.12}$$

for any $0 < \delta \leq 1$ *such that*

$$\max\{\Delta(S'[i,j], R'[i,j]),\ \Delta(S''[i,j], R''[i,j])\} \geq \delta. \tag{4.4.13}$$

Proof. By hypothesis, $\Delta(S'[i,j], R'[i,j]) \geq \delta$ or $\Delta(S''[i,j], R''[i,j]) \geq \delta$. Without loss of generality, assume the former. Abbreviating $J = [i,j]$, we have

$$\frac{|D'|_J + |I'|_J}{|D'|_J + |G'|_J} \geq \delta,$$

which along with $\delta > 0$ gives

$$|D'|_J + |G'|_J \leq \frac{|D'|_J + |I'|_J}{\delta}. \tag{4.4.14}$$

Now

$$\begin{aligned}
\frac{\tilde{\|}J\tilde{\|}-1}{2} &\leq |G' \cup D' \cup I'|_J \\
&= |G'|_J + |D' \cup I'|_J \\
&= |G'|_J + |D'|_J + |I' \setminus D'|_J \\
&\leq \frac{|D'|_J + |I'|_J}{\delta} + |I' \setminus D'|_J \\
&\leq \frac{|D'|_J + |I'|_J + |I' \setminus D'|_J}{\delta} \\
&= \frac{|I'|_J + |I' \cup D'|_J}{\delta} \\
&\leq \frac{|I|_J + |I \cup D|_J}{\delta} \\
&= \frac{2|D|_J}{\delta},
\end{aligned}$$

first step follows from Lemma 4.26, the second uses Lemma 4.21 (vii), (viii), the fourth is valid by (4.4.14), the fifth uses $0 < \delta \leq 1$, and the last step is immediate from Lemma 4.21 (iv). □

Finally, we derive a useful bound on the virtual length of an interval in terms of the number of excellent and bad transmissions in it.

Lemma 4.29. *For any interval J,*

$$\tilde{\|}J\tilde{\|} \leq 2(|B|_J + |E'|_J) + 1, \tag{4.4.15}$$

$$\tilde{\|}J\tilde{\|} \leq 2(|B|_J + |E''|_J) + 1. \tag{4.4.16}$$

Proof. By symmetry, it suffices to prove (4.4.15). We have

$$\begin{aligned} D' \setminus I' &= I'' \setminus D'' \\ &\subseteq I'' \\ &\subseteq I'' \cup (G'' \setminus E'') \\ &= (I'' \cup G'') \setminus E'' \\ &= B'', \end{aligned} \tag{4.4.17}$$

where the first and fourth steps use parts (vi) and (vii), respectively, of Lemma 4.21. Now (4.4.15) can be verified as follows:

$$\begin{aligned} \frac{\tilde{\|}J\tilde{\|} - 1}{2} &\leq |G' \cup I' \cup D'|_J \\ &= |G' \cup I'|_J + |D' \setminus (G' \cup I')|_J \\ &= |G' \cup I'|_J + |D' \setminus I'|_J \\ &= |E'|_J + |(G' \cup I') \setminus E'|_J + |D' \setminus I'|_J \\ &= |E'|_J + |B'|_J + |D' \setminus I'|_J \\ &\leq |E'|_J + |B'|_J + |B''|_J \\ &= |E'|_J + |B|_J, \end{aligned}$$

where the first step is valid by Lemma 4.4.8, the third step uses Lemma 4.21 (viii), the fourth step follows from the containment $E' \subseteq G'$, the fifth step applies the definition of B', the sixth step is immediate from (4.4.17), and the final step is justified by Lemma 4.24. □

4.4.6. Virtual corruption rate. In keeping with our focus on events rather than transmissions, we define

$$\operatorname{corr} J = \frac{|D \cap J|}{\tilde{\|} J \tilde{\|}}$$

for any interval J. We refer to this quantity as the *virtual corruption rate* of J. The next lemma shows that over the course of the execution, the virtual corruption rate is relatively low.

Lemma 4.30. *Assumptions* (i) *and* (ii) *in Theorem* 4.19 *imply*

$$\operatorname{corr}[1, T] \leq \frac{1}{4} - \varepsilon \tag{4.4.18}$$

and

$$\operatorname{corr}[1, T] \leq \frac{1}{4} - \frac{\varepsilon}{2}, \tag{4.4.19}$$

respectively.

Proof. Assumption (i) states that the total number of attacks does not exceed a $\frac{1}{4} - \varepsilon$ fraction of the worst-case communication cost of the interactive coding scheme. Formulaically,

$$|D| \leq \left(\frac{1}{4} - \varepsilon\right) \cdot 2N.$$

As a result,

$$\operatorname{corr}[1, T] = \frac{|D|}{\tilde{\|}[1, T]\tilde{\|}} \leq \frac{1}{4} - \varepsilon,$$

where the second step uses Lemma 4.27.

Assumption (ii) states that

$$T_{\mathrm{subs}} + \frac{3}{4} T_{\mathrm{oos}} \leq \left(\frac{1}{4} - \varepsilon\right) T,$$

where T_{subs} and T_{oos} denote the total number of substitution attacks and the total number of out-of-sync attacks, respectively. Straightforward manipulations now reveal that

$$\frac{T_{\mathrm{subs}} + T_{\mathrm{oos}}}{T + T_{\mathrm{oos}}} \leq \frac{1}{4} - \frac{\varepsilon}{2}.$$

By definition,

$$|D| = T_{\mathrm{subs}} + T_{\mathrm{oos}}.$$

On the other hand, the defining equation (4.4.7) of virtual length reveals that $\tilde{\|}Z\tilde{\|}$ for any set Z is the total number of transmissions in Z plus the total number of out-of-sync attacks in Z. In particular,

$$\tilde{\|}[1, T]\tilde{\|} = T + T_{\mathrm{oos}}.$$

The last three equations immediately give (4.4.19). □

4.4.7. Finish times. Let $e_1, e_2, \ldots, e_n$ be the edges of the unique root-to-leaf path in $X \cup Y$, listed in increasing order of depth. For $i = 1, 2, \ldots, n$, define f_i to be the index of the first transmission when e_i is sent (whether or not that transmission is subject to an attack). If e_i is never sent, we define $f_i = \infty$. For notational convenience, we also define $f_0 = f_{-1} = f_{-2} = \cdots = 0$. Recall from the description of the interactive coding scheme that Alice never sends an edge e unless she has previously sent all proper predecessors of e in X, and analogously for Bob. This gives

$$f_1 \le f_3 \le f_5 \le \cdots,$$
$$f_2 \le f_4 \le f_6 \le \cdots.$$

The overall sequence $f_1, f_2, f_3, f_4, f_5, f_6, \ldots$ need not be in sorted order, however, due to interference by the adversary. We abbreviate

$$\overline{f_i} = \max\{0, f_1, f_2, \ldots, f_i\}.$$

By basic arithmetic,

$$[\overline{f_{i-1}}, \overline{f_i}) = [\overline{f_{i-1}}, f_i), \qquad i = 1, 2, \ldots, n. \tag{4.4.20}$$

We now bound the virtual length of any such interval in terms of the number of bad transmissions in it, thereby showing that Alice and Bob make rapid progress as long as they do not experience too many attacks.

Lemma 4.31. *For any integers i and t with $\overline{f_{i-1}} \le t < f_i$,*

$$\tilde{\|}[\overline{f_{i-1}}, t]\tilde{\|} \le 2|B|_{[\overline{f_{i-1}}, t]} + 3. \tag{4.4.21}$$

Proof. We will only treat the case of i odd; the proof for even i can be obtained by swapping the roles of Alice and Bob below.

Consider any transmission $j \in E' \cap [\overline{f_{i-1}}, f_i)$. Lemma 4.23 ensures that on receipt of transmission j, Alice is able to correctly recover the set of edges that Bob has sent her by that time, which includes $e_2, e_4, e_6, \ldots, e_{i-1}$. At that same time, Alice has sent Bob $e_1, e_3, e_5, \ldots, e_{i-2}$ but not e_i, as one can verify from $j \in [\overline{f_{i-1}}, f_i)$. Therefore, the arrival of transmission j causes Alice either to exit or to immediately send e_i. Either way, the interval $[\overline{f_{i-1}}, f_i)$ does not contain any transmissions numbered $j+1$ or higher. We conclude that there is at most one transmission in $E' \cap [\overline{f_{i-1}}, f_i)$, and in particular

$$|E'|_{[\overline{f_{i-1}},t]} \leq 1.$$

This upper bound directly implies (4.4.21) in light of Lemma 4.29. □

4.4.8. The progress lemma. We have reached the technical centerpiece of our analysis. The result that we are about to prove shows that any sufficiently long execution of the interactive coding scheme with a sufficiently low virtual corruption rate allows Alice and Bob to exchange all the n edges of the unique root-to-leaf path in $X \cup Y$, and moreover this progress is not "undone" by any subsequent attacks by the adversary. The proof uses amortized analysis in an essential way.

LEMMA 4.32 (Progress lemma). *Let $t \in \{1, 2, \ldots, T\}$ be given with*

$$\tilde{\|}[1,t]\tilde{\|} \geq \frac{n+2}{\alpha}, \tag{4.4.22}$$

$$\operatorname{corr}[1,t] \leq \frac{1}{4} - \alpha. \tag{4.4.23}$$

Then there is an integer $t^* \leq t$ such that

$$[\overline{f_n}, t^*) \cap E' \neq \varnothing, \tag{4.4.24}$$

$$[\overline{f_n}, t^*) \cap E'' \neq \varnothing, \tag{4.4.25}$$

$$\Delta(S'[i,t], R'[i,t]) < 1 - \alpha, \qquad i = 1, 2, \ldots, t^*, \tag{4.4.26}$$

$$\Delta(S''[i,t], R''[i,t]) < 1 - \alpha, \qquad i = 1, 2, \ldots, t^*. \tag{4.4.27}$$

Proof. Equations (4.4.26) and (4.4.27) hold vacuously for $t^* = 0$. In what follows, we will take $t^* \in \{0, 1, 2, \ldots, t\}$ to be the *largest* integer for which (4.4.26) and (4.4.27) hold. For the sake of contradiction, assume that at least one of the remaining desiderata (4.4.24), (4.4.25) is violated, whence

$$\tilde{\|}[\overline{f_n}, t^*)\tilde{\|} \leq 2|B|_{[\overline{f_n},t^*)} + 1 \tag{4.4.28}$$

by Lemma 4.29. The proof strategy is to show that (4.4.28) is inconsistent with the hypothesis of the lemma. To this end, let $n^* \in \{0, 1, 2, \ldots, n\}$ be the largest integer such that $\overline{f_{n^*}} \leq t^*$. Then we have the partition

$$[0, t] = [\overline{f_0}, \overline{f_1}) \cup [\overline{f_1}, \overline{f_2}) \cup \cdots \cup [\overline{f_{n^*-1}}, \overline{f_{n^*}}) \cup [\overline{f_{n^*}}, t^*) \cup \{t^*\} \cup (t^*, t].$$

The bulk of our proof is concerned with bounding the virtual length of each of the intervals on the right-hand side.

To begin with,

$$\begin{aligned} \tilde{\|}[\overline{f_{i-1}}, \overline{f_i})\tilde{\|} &= \tilde{\|}[\overline{f_{i-1}}, f_i)\tilde{\|} \\ &\leq 2|B|_{[\overline{f_{i-1}},f_i)} + 3 \\ &\leq 2|B|_{[\overline{f_{i-1}},\overline{f_i})} + 3 \end{aligned} \tag{4.4.29}$$

for any $i = 1, 2, \ldots, n^*$, where the first and third steps use (4.4.20), and the second step follows from Lemma 4.31. Next, the upper bound

$$\tilde{\|}[\overline{f_{n^*}}, t^*)\tilde{\|} \leq 2|B|_{[\overline{f_{n^*}},t^*)} + 3 \tag{4.4.30}$$

follows from Lemma 4.31 if $n^* < n$ and from (4.4.28) if $n^* = n$. The virtual length of the singleton interval $\{t^*\}$ can be bounded from first principles:

$$\tilde{\|}\{t^*\}\tilde{\|} \leq 2. \tag{4.4.31}$$

Finally, recall from the definition of t^* that either $\max\{\Delta(S'[t^*+1,t], R'[t^*+1,t]), \Delta(S''[t^*+1,t], R''[t^*+1,t])\} \geq 1-\alpha$ or $t^* = t$, leading to

$$\tilde{\|}(t^*, t]\tilde{\|} \leq \frac{4}{1-\alpha}|D|_{(t^*,t]} + 1 \tag{4.4.32}$$

by Lemma 4.28 in the former case and trivially in the latter.

Putting everything together, we obtain

$$\begin{aligned}
\tilde{\|}[1, t]\tilde{\|} &\leq 2|B|_{[0,t^*)} + 3(n^*+1) + 2 + \frac{4}{1-\alpha}|D|_{(t^*,t]} + 1 \\
&\leq \frac{4}{1-\alpha}|D|_{[0,t^*)} + 3(n^*+1) + 2 + \frac{4}{1-\alpha}|D|_{(t^*,t]} + 1 \\
&\leq \frac{4}{1-\alpha}|D|_{[0,t]} + 3n + 6 \\
&\leq \frac{4}{1-\alpha}|D|_{[0,t]} + 3\alpha\tilde{\|}[1, t]\tilde{\|}, \qquad (4.4.33)
\end{aligned}$$

where the first step is the result of adding (4.4.29)–(4.4.32), the second step applies Lemma 4.25, and the final step uses (4.4.22). Since $0 < \alpha < 1$, the conclusion of (4.4.33) is equivalent to

$$\operatorname{corr}[1, t] \geq \frac{(1-3\alpha)(1-\alpha)}{4},$$

which is inconsistent with (4.4.23). We have obtained the desired contradiction and thereby proved that t^* satisfies each of the properties (4.4.24)–(4.4.27). □

4.4.9. Finishing the proof. We have reached a "master theorem," which gives a sufficient condition for Alice and Bob to assign the correct value to their corresponding copies of the *out* variable. Once established, this result will allow us to easily finish the proof of Theorem 4.19.

THEOREM 4.33. *Consider a point in time when Alice updates her out variable, and fix a corresponding integer $t \leq T$ such that $\not*(R'[1, t])$ is the complete sequence of* *symbols that Alice has received by that time. Assume that*

$$\tilde{\|}[1,t]\tilde{\|} \geq \frac{n+2}{\alpha}, \tag{4.4.34}$$

$$\operatorname{corr}[1,t] \leq \frac{1}{4} - \alpha. \tag{4.4.35}$$

Then as a result of the update, out is assigned the leaf vertex in the unique root-to-leaf path in $X \cup Y$. An analogous theorem holds for Bob.

Observe that Theorem 4.33 makes no assumption as to the actual timing of the update to *out*. It may happen that the update takes place in response to the t^{th} transmission; but it may also take place significantly earlier, due to out-of-sync attacks.

Proof of Theorem 4.33. We will only prove the claim for Alice; the proof of Bob is entirely analogous. Lemma 4.32 implies the existence of $j' \in E'$ and $j'' \in E''$ such that

$$\overline{f_n} \leq j' < t, \tag{4.4.36}$$

$$\overline{f_n} \leq j'' < t, \tag{4.4.37}$$

$$\Delta(S'[j'+1,t], R'[j'+1,t]) < 1-\alpha, \tag{4.4.38}$$

$$\Delta(S''[j''+1,t], R''[j''+1,t]) < 1-\alpha. \tag{4.4.39}$$

By the definition of E' and E'',

$$\operatorname{SD}(S'[1,j'], R'[1,j']) < 1-\alpha, \tag{4.4.40}$$

$$\operatorname{SD}(S''[1,j''], R''[1,j'']) < 1-\alpha. \tag{4.4.41}$$

As a result,

$$\begin{aligned}
&\mathrm{SD}_{\overline{\ast}(S'[1,j'])}(S'[1,t], R'[1,t]) \\
&\quad = \mathrm{SD}_{\overline{\ast}(S'[1,j'])}(S'[1,j']S'[j'+1,t],\ R'[1,j']R'[j'+1,t]) \\
&\quad \leq \max\{\mathrm{SD}_{\overline{\ast}(S'[1,j'])}(S'[1,j'], R'[1,j']),\ \Delta(S'[j'+1,t], R'[j'+1,t])\} \\
&\quad \leq \max\{\mathrm{SD}(S'[1,j'], R'[1,j']),\ \Delta(S'[j'+1,t], R'[j'+1,t])\} \\
&\quad < 1-\alpha, \qquad\qquad (4.4.42)
\end{aligned}$$

where the second step is valid by Proposition 4.12 (iv), the third step uses (4.2.5), and the final step is immediate from (4.4.38) and (4.4.40).

When Alice updates her *out* variable, the sequence of symbols that she has received is $\not\ast(R'[1,t])$. By (4.4.42) and Theorem 4.17,

$$\begin{aligned}
\textsc{Decode}_{C,\alpha}(\not\ast(R'[1,t])) &\succeq (\not\ast(S'[1,t]))_{\leq \overline{\ast}(S'[1,j'])} \\
&= \not\ast(S'[1,j']).
\end{aligned}$$

Therefore, just prior to updating *out*, Alice is able to correctly recover the prefix $\not\ast(S'[1,j'])$ of the sequence of symbols sent to her by Bob. By Lemma 4.22, this means that she correctly recovers the complete set of edges encoded by the string $\not\ast(S'[1,j'])$. By (4.4.36), this prefix $\not\ast(S'[1,j'])$ contains the encoding of every edge of Y that appears in the root-to-leaf path in $X \cup Y$. Moreover, every edge encoded in $\not\ast(S'[1,j'])$ is correct in that it is an element of Y. Alice's pseudocode now ensures that she assigns to *out* the leaf vertex on the unique root-to-leaf path in $X \cup Y$.

The proof for Bob is entirely analogous, with (4.4.37), (4.4.39), (4.4.41), and j'' playing the role of (4.4.36), (4.4.38), (4.4.40), and j', respectively. □

We are now in a position to establish the main result of this section.

Proof of Theorem 4.19. Recall that $n = |\pi|$ denotes the communication cost of the original protocol, and $\varepsilon > 0$ is a constant in the statement of Theorem 4.19. Consider the interactive coding scheme given by Algorithms 2 and 3, with parameters set according to

$$\alpha = \frac{\varepsilon}{4}, \tag{4.4.43}$$

$$N = \left\lceil \frac{n+4}{2\alpha} \right\rceil. \tag{4.4.44}$$

By (4.4.2), the coding scheme uses an alphabet of size at most $(|\Sigma| \cdot n/\varepsilon)^{O(1/\varepsilon)} = O(|\Sigma| \cdot n)^{O(1)} = O(|\Sigma| \cdot |\pi|)^{O(1)}$. Furthermore, by (4.4.3), the combined number of transmissions sent by Alice and Bob does not exceed $2N = O(n) = O(|\pi|)$.

It remains to show that when the communication stops, *out* is set for both Alice and Bob to the leaf vertex on the unique root-to-leaf path in $X \cup Y$. To this end, note from (4.4.43) and Lemma 4.30 that

$$\operatorname{corr}[1, T] \leq \frac{1}{4} - 2\alpha. \tag{4.4.45}$$

By (4.4.44) and Lemma 4.27,

$$\tilde{\|}[1, T]\tilde{\|} > \frac{n+4}{\alpha} \tag{4.4.46}$$

and therefore

$$\tilde{\|}[1, T-1]\tilde{\|} > \frac{n+2}{\alpha}. \tag{4.4.47}$$

Also,

$$\begin{aligned}
\operatorname{corr}[1,T-1] &\leq \frac{\tilde{\|}[1,T]\tilde{\|}}{\tilde{\|}[1,T-1]\tilde{\|}} \cdot \operatorname{corr}[1,T] \\
&\leq \left(1 + \frac{2}{\tilde{\|}[1,T-1]\tilde{\|}}\right) \cdot \operatorname{corr}[1,T] \\
&\leq \left(1 + \frac{2\alpha}{n+2}\right) \cdot \left(\frac{1}{4} - 2\alpha\right) \\
&\leq \frac{1}{4} - \alpha, \qquad (4.4.48)
\end{aligned}$$

where the third step uses (4.4.45) and (4.4.47). Now, consider the last time that Alice and Bob update their copies of *out*. The complete sequence of symbols that Alice has received at the time of her last update is $\not\ast(R'[1,T-1])$ or $\not\ast(R'[1,T])$. Likewise, the complete sequence of symbols that Bob has received at the time of his last update is $\not\ast(R''[1,T-1])$ or $\not\ast(R''[1,T])$. By (4.4.45)–(4.4.48) and Theorem 4.33, both players set *out* to the leaf vertex in the unique root-to-leaf path in $X \cup Y$. This completes the proof of Theorem 4.19. □

4.5. A coding scheme with a constant-size alphabet

In this section, we will adapt the proof of Theorem 4.19 to use an alphabet of constant size. This modification will yield the main result of this chapter (Theorems 4.1 and 4.2), which we restate here for the reader's convenience.

THEOREM 4.34. *Fix an arbitrary constant $\varepsilon > 0$, and let π be an arbitrary protocol with alphabet Σ. Then there exists an interactive coding scheme for π with alphabet size $O(1)$ and communication cost $O(|\pi| \log |\Sigma|)$ that tolerates*

(i) *corruption rate $\frac{1}{4} - \varepsilon$;*

(ii) *any normalized corruption rate $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$ with $\varepsilon_{\text{subs}} + \frac{3}{4}\varepsilon_{\text{oos}} \leq \frac{1}{4} - \varepsilon$.*

In Section 4.5.8, we further generalize Theorem 4.34(i) to the setting when Alice and Bob need to be ready with their answers by a certain round (based on each player's own counting) rather than when the communication stops. In that setting, too, our interactive coding scheme is optimal and matches the lower bound due to Braverman et al. [**25**]. At a high level, our proofs of Theorem 4.34 and its generalization are similar to the proof of Theorem 4.19 in the previous section, and we will be able to reuse most of the auxiliary machinery developed there. The principal point of departure is a new way of encoding and transferring edges, which in turn requires subtle modifications to the amortized analysis.

4.5.1. Edge representation and transfer. We may assume without loss of generality that π is in canonical form, which can be achieved for any protocol at the expense of doubling its communication cost. Canonical form allows us to identify Alice's input with a set X of odd-depth edges of the protocol tree for π, and Bob's input with a set Y of even-depth edges. Execution of π corresponds to following the unique root-to-leaf path in $X \cup Y$, whose length we denote by

$$n = |\pi|.$$

Recall that our previous interactive coding scheme in Section 4.4 involved Alice and Bob sending each other edges from their respective input sets X and Y, with each transmission representing precisely one such edge. The new coding scheme also amounts to Alice and Bob exchanging edges from their respective input sets. This time, however, any given transmission will contain information about as many as Λ^2 edges, where $\Lambda = \Lambda(\varepsilon) > 0$ is a constant to be chosen later. Moreover, to accommodate the size restriction on the alphabet, the encoding of any given edge will now be split across multiple transmissions. We say that a transmission *fulfills* an edge e if it carries the last bit of e's encoding.

Our approach to the encoding and transfer of edges is inspired by the interactive coding schemes with constant-size alphabets due to Braverman and Rao [**26**] and Braverman et al. [**25**]. We adapt their encoding and transfer in several ways to support our more general setting and to make the overall proof simpler. A detailed technical exposition follows.

Edge encoding. We will keep the policy that Alice does not start sending an edge e unless she has already fulfilled all predecessors of e in X, and likewise for Bob. This makes it possible for the sender to encode an edge e by referring to the previously transmitted grandparent of e. Specifically, an edge is now encoded as a triple (m, j, σ), where m is the number of transmissions sent by the sender since his or her most recent transmission that fulfilled the grandparent of e; the index $j \in \{1, 2, 3, \ldots, \Lambda^2\}$ identifies that grandparent among the up to Λ^2 edges featured in that transmission; and $\sigma \in \Sigma \times \Sigma$ identifies e relative to that grandparent. As a base case, an edge of depth 1 or 2 is encoded by a triple (m, j, σ) where m is the number of transmissions sent by the sender since the beginning of time, and j is ignored. Note that how an edge is encoded is highly context-sensitive in that it depends on previous transmissions by the sender. As a result, whenever we speak of *the* encoding of an edge e, we are referring to the encoding of e at a particular time that will be clear from the context.

Chunking. A constant-size alphabet makes it in general impossible to deliver the entire encoding of an edge in a single transmission. Instead, we split the encoding of every edge into *chunks*. A chunk contains a single bit of the encoding of the edge as well as 3 bits of metadata. Thus, the number of chunks needed to transfer an edge is equal to the bit length of e's encoding. Alice and Bob each maintain data structures called *encoding* and *numBitsSent*, indexed by edges. These data structures store, for each of the edges currently being transferred, its encoding and the number of bits sent so far.

Parallelism. Rather than send edges one by one, each player will send up to Λ^2 edges in parallel. To see the intuitive reason for doing so, consider the transfer of a typical edge e, which spans multiple transmissions. As Alice sends e chunk by chunk to Bob, she simultaneously receives information from him, which in turn may lead her to believe that she should be sending an edge other than e. The problem is, she can never be sure! Simply aborting the transfer of e is wasteful if e later turns out to be the right edge to send. Instead, we allow transfer of several edges in parallel and use an additional, credit-based mechanism for identifying and aborting unpromising transfers.

Specifically, each player maintains an ordered list L of edges that he or she is currently transferring. New edges are inserted in L at the front rather than back, reflecting that view that new information should be prioritized over old. To prepare a transmission, a player looks at the first Λ^2 edges in L and takes a chunk of each. If L has fewer than Λ^2 edges, the player simply takes a chunk of each edge in L. The concatenation of these chunks, ordered the same way as the corresponding edges in L, forms a *page*, which we view as a symbol from an auxiliary alphabet Σ_{in}. Since an edge chunk is a 4 bits long, the size of Σ_{in} is bounded by a constant:

$$|\Sigma_{\text{in}}| = \sum_{i=0}^{\Lambda^2} 2^{4i} = \frac{16^{\Lambda^2+1} - 1}{15}. \tag{4.5.1}$$

Credit. As a crucial component of the transfer scheme, Alice and Bob each maintain a data structure called *credit*. This data structure is indexed by edges and stores the amount of "funds" available to pay for the transfer of any given edge e. The credit of every edge is initialized to 0 at the beginning, and remains nonnegative from then on. Every receive-send cycle identifies an edge e to send, which then gets a credit increase of Λ and is additionally inserted in L unless it is already there. Any time an edge chunk is sent, the credit of the corresponding edge is decreased by 1. An edge remains in L until its credit reaches 0 or until its last chunk is sent, whichever comes first. At that point, the edge is removed from L.

Metadata. The purpose of the metadata in each edge chunk is to allow the receiver to correctly piece together the encodings of the edges. A chunk for an edge e is always prepared at send time rather than in advance and includes the following four bits: the next bit of the encoding of e; a bit to indicate if this is the first chunk for e; a bit to indicate if this is the last chunk for e; and a bit to indicate if e's credit has reached zero. The last two bits alert the receiver to the removal of e from the sender's edge list.

4.5.2. The simulation. Algorithm 4 gives the pseudocode to support our edge encoding and transfer scheme. The pseudocode is identical for Alice and Bob and features the following three operations.

(i) AddEdge is executed once by each player during his or her receive-send cycle. As an argument, it receives an edge which that player wants to send next. If e is already on the player's edge list, AddEdge simply increments e's credit by Λ. If not, AddEdge increments e's credit by Λ, computes an encoding of e relative to the player's current transmission count, and adds e to the edge list ahead of any existing edges.

(ii) NextChunk receives as an argument an edge e and returns the next 4-bit chunk of that edge, based on the stored encoding of e and the number of bits of e's encoding sent so far. This procedure uses *numBitsSent*(e), *credit*(e), and *encoding*(e) to correctly set the metadata for the chunk. It then updates *numBitsSent*(e) and *credit*(e) to reflect the remaining number of bits to send and the edge's available credit.

```
Global variables: $encoding, numBitsSent, credit, L$

Procedure AddEdge($e, i$)
    $credit(e) \leftarrow credit(e) + \Lambda$
    if $e \notin L$ then
        $encoding(e) \leftarrow$ encoding of $e$ based on current transmission count $i$
        $numBitsSent(e) \leftarrow 0$
        prepend $e$ to $L$, ahead of any existing edges
    end

Procedure NextChunk($e$)
                                                    // Update edge statistics
    $numBitsSent(e) \leftarrow numBitsSent(e) + 1$
    $credit(e) \leftarrow credit(e) - 1$
                                                    // Compute edge chunk
    return $(encoding(e))_{numBitsSent(e)}$
            $\circ\, \mathbf{I}[numBitsSent(e) = 1]$
            $\circ\, \mathbf{I}[numBitsSent(e) = |encoding(e)|]$
            $\circ\, \mathbf{I}[credit(e) = 0]$

Procedure NextPage()
    $page \leftarrow$ NextChunk($L[1]$)
            $\circ$ NextChunk($L[2]$)
            $\circ \cdots$
            $\circ$ NextChunk($L[\min\{\Lambda^2, |L|\}]$)
                                                    // Clean up the edge list
    foreach $e \in L$ do
        if $credit(e) = 0$ or $numBitsSent(e) = |encoding(e)|$ then
            remove $e$ from $L$
        end
    end
    return $page$
```

Algorithm 4: Edge operations (identical for Alice and Bob). In the pseudocode above, $\circ$ denotes string concatenation, $|L|$ denotes the number of edges in L, and $L[i]$ denotes the i^{th} edge in L.

(iii) NextPage is the procedure that assembles the next page to send. The page is made up of at most Λ^2 chunks, one for each of the first Λ^2 edges on the

```
Input: $X$ (set of Alice's edges)
$L \leftarrow \varnothing$
$credit(e) \leftarrow 0$ for every edge $e$
$e \leftarrow$ the edge in $X$ incident to the root
AddEdge$(e, 1)$
$page \leftarrow$ NextPage()
encode and send $page$
foreach $i = 1, 2, 3, \ldots, N$ do
    receive a symbol $r_i \in \Sigma_{\text{out}}$
    $s \leftarrow \text{Decode}_{C,\alpha}(r_1 r_2 \ldots r_i)$
    interpret $s$ as a sequence $B$ of even-depth edges
    $\ell \leftarrow$ maximum length of a rooted path in $X \cup B$
    compute the shortest prefix of $B$ s.t. $X \cup B$ contains a rooted path of length $\ell$,
      and let $P$ be the rooted path so obtained
    $out \leftarrow$ deepest vertex in $P$
    if $i \leq N - 1$ then
        $e \leftarrow$ the deepest edge in $P \cap X$ whose proper predecessors in $X$ have all
          been sent
        AddEdge$(e, i + 1)$
        $page \leftarrow$ NextPage()
        encode and send $page$
    end
end
```

Algorithm 5: Coding scheme for Alice

edge list. The chunks are prepared using NextChunk. Once the page is assembled, NextPage updates the edge list by removing edges that have been fully sent or have no credit left.

The overall interactive coding scheme is given by Algorithms 5 and 6 for Alice and Bob, respectively. The main novelty relative to the scheme of Section 4.4 are the calls to AddEdge and NextPage, which a player executes as soon as he or she has

```
Input: $Y$ (set of Bob's edges)
$L \leftarrow \varnothing$
$\mathit{credit}(e) \leftarrow 0$ for every edge $e$
foreach $i = 1, 2, 3, \ldots, N$ do
    receive a symbol $r_i \in \Sigma_{\text{out}}$
    $s \leftarrow \textsc{Decode}_{C,\alpha}(r_1 r_2 \ldots r_i)$
    interpret $s$ as a sequence $A$ of odd-depth edges
    $\ell \leftarrow$ maximum length of a rooted path in $Y \cup A$
    compute the shortest prefix of $A$ s.t. $Y \cup A$ contains a rooted path of length $\ell$,
      and let $P$ be the rooted path so obtained
    $\mathit{out} \leftarrow$ deepest vertex in $P$
    $e \leftarrow$ the deepest edge in $P \cap Y$ whose proper predecessors in $Y$ have all been
      sent
    $\textsc{AddEdge}(e, i)$
    $\mathit{page} \leftarrow \textsc{NextPage}()$
    encode and send $\mathit{page}$
end
```

Algorithm 6: Coding scheme for Bob

identified an edge e to send. Apart from that, the remarks made in Section 4.4 apply here in full. In particular, $\alpha = \alpha(\varepsilon) \in (0,1)$ and $N = N(n, \alpha)$ are parameters to be chosen later. We set

$$\Lambda = \left\lceil \frac{2}{\alpha} \right\rceil \tag{4.5.2}$$

and fix an arbitrary α-good code $C\colon \Sigma_{\text{in}}^* \to \Sigma_{\text{out}}^*$ whose existence is ensured by Theorem 4.9. That theorem implies, in view of (4.5.1) and (4.5.2), that

$$|\Sigma_{\text{out}}| \leq 2^{O(1/\alpha^3)}. \tag{4.5.3}$$

Alice and Bob use C to encode every transmission. In particular, the encoded symbol from Σ_{out} at any given point depends not only on the symbol from Σ_{in} being transmitted but also on the content of the previous transmissions by the sender. The decoding is again done using the $\textsc{Decode}_{C,\alpha}$ algorithm of Theorem 4.17. Note from the pseudocode that Alice and Bob send at most N transmissions each.

It remains to elaborate on the decoding and interpretation steps in the interactive coding scheme. To do so, we first prove that the sequence of pages sent by one of the players at any given point reveals the sequence of edges that that player has fulfilled so far.

LEMMA 4.35. *Consider an arbitrary point in time, and let $p_1, p_2, \ldots, p_t \in \Sigma_{\text{in}}$ be the sequence of pages sent by one of the players so far. Then that sequence uniquely identifies the corresponding sequence of edges $e_1, e_2, \ldots, e_{t'}$ fulfilled by that player.*

Proof. We first reconstruct as completely as possible the sender's state at the times when each of the pages $p_1, p_2, \ldots, p_t$ has just been assembled. Specifically, we determine the length of the sender's edge list, the transmission status of every edge on the edge list (in progress, aborted, or fulfilled), and the corresponding part of the encoding transferred for every edge so far. This reconstruction process involves working inductively through the page sequence $p_1, p_2, \ldots, p_t$ and using the metadata to identify when an edge is new, in progress, aborted, or fulfilled. Recall that there is at most one new edge per page, and it is always inserted at the *front* of the edge list. The first stage reconstructs the complete list of edge encodings sent so far by the sender, along with the final status of each encoding (in progress, aborted, or fulfilled), and the start and end times of each fulfilled encoding. We then interpret the fulfilled encodings as a sequence $(m_1, j_1, \sigma_1), (m_2, j_2, \sigma_2), \ldots, (m_{t'}, j_{t'}, \sigma_{t'})$ of edge representations. Using the end times of the fulfilled encodings and their indices inside the pages than fulfilled them, we can reconstruct the corresponding sequence of edges $e_1, e_2, \ldots, e_{t'}$ via an inductive process analogous to that in Lemma 4.20. □

With Lemma 4.35 in hand, the decoding and interpretation steps in lines 9–10 for Alice and lines 5–6 for Bob are implemented the same way they were for a large alphabet. Specifically, the decoding step produces a codeword s of C, which by Fact 4.8 corresponds to a unique string in Σ_{in}^*. This string is by definition a sequence of pages $p_1, p_2, p_3, \ldots$, from which the receiving party can reconstruct the corresponding sequence of fulfilled edges using the inductive procedure of Lemma 4.35. It may happen that the page sequence $p_1, p_2, p_3, \ldots$ is syntactically malformed; in that case, the receiving party interrupts the interpretation process at the longest prefix of $p_1, p_2, p_3, \ldots$ that corresponds to a legitimate sequence of edges. This completes the interpretation step, yielding a sequence of edges A for Bob and B for Alice.

Analogous to the interactive coding scheme of Section 4.4, Alice and Bob each maintain a variable called *out*. In Sections 4.5.3–4.5.7 below, we will examine an arbitrary but fixed execution of the interactive coding scheme. In particular, we will henceforth consider the inputs X and Y and the adversary's actions to be fixed. We allow any behavior by the adversary as long as it meets one of the criteria (i), (ii) in Theorem 4.34. We will show that as soon as the communication stops, *out* is set for both Alice and Bob to the leaf vertex of the unique root-to-leaf path in $X \cup Y$. This will prove Theorem 4.34.

4.5.3. Fundamental notions and facts. We adopt the notation and definitions of Sections 4.4.2–4.4.6 in their entirety. These items carry over without any changes because they pertain to the lowest level of abstraction (the "data link layer," as it were), which cannot distinguish between the old and new interactive coding schemes. As a consequence, all results proved in Sections 4.4.2–4.4.6 apply here in full, with the exception are Lemmas 4.22 and 4.23 whose wording needs to be clarified by replacing "sent edges" with "fulfilled edges." The result of this cosmetic modification is as follows.

Lemma 4.36. *Let $t \in \{1, 2, \ldots, T\}$ be given. Then:*

(i) *the string ${\not*}(S'[1,t])$ uniquely identifies the sequence of protocol tree edges that Bob fulfills over the course of transmissions $1, 2, \ldots, t$;*

(ii) *the string ${\not*}(S''[1,t])$ uniquely identifies the sequence of protocol tree edges that Alice fulfills over the course of transmissions $1, 2, \ldots, t$.*

Proof. By symmetry, it suffices to prove the former claim. By Fact 4.8, the codeword ${\not*}(S'[1,t]) \in \Sigma_{\mathrm{out}}^*$ corresponds to a unique string in Σ_{in}^*, which is the sequence of pages that Bob sends Alice over the course of the first t transmissions. By Lemma 4.35, this sequence of pages uniquely identifies the corresponding fulfilled edges. □

Lemma 4.37. *Let $t \in \{1, 2, \ldots, T\}$ be given.*

(i) *If $t \in E'$, then on receipt of transmission t, Alice is able to correctly recover the complete sequence of edges that Bob has fulfilled by that time.*

(ii) *If $t \in E''$, then on receipt of transmission t, Bob is able to correctly recover the complete sequence of edges that Alice has fulfilled by that time.*

Proof. By symmetry, it again suffices to prove the former claim. Let $t \in E'$. Then by definition, $\mathrm{SD}(S'[1,t], R'[1,t]) < 1 - \alpha$. Taking $k = \infty$ in Theorem 4.17, we conclude that $\textsc{Decode}_{C,\alpha}({\not*}(R'[1,t])) = {\not*}(S'[1,t])$. This means that on receipt of transmission t, Alice is able to correctly recover the entire codeword ${\not*}(S'[1,t])$ that Bob has sent her so far. By Lemma 4.36, this in turn makes it possible for Alice to correctly identify the corresponding sequence of fulfilled edges. □

4.5.4. Full pages. Recall that a page can contain at most Λ^2 edge chunks. If a page contains exactly Λ^2 chunks, we call it *full.* We define $F' \subseteq \{1, 2, \ldots, T\}$ as the set of transmissions where Alice sends a full page, and analogously $F'' \subseteq \{1, 2, \ldots, T\}$

as the set of transmissions where Bob sends a full page. In other words,

$$F' = \{i : S''[i,i] \text{ is a full page}\},$$

$$F'' = \{i : S'[i,i] \text{ is a full page}\}.$$

We abbreviate

$$F = F' \cup F''.$$

The following lemma, due to Braverman et al. [**25**, Lemma D.1], shows that full pages are relatively uncommon.

PROPOSITION 4.38 (Braverman et al.). *For any interval J such that $1 \in J$,*

$$|F|_J \leq \frac{|J \cap \{1,2,3,\ldots,T\}|}{\Lambda}.$$

Proof (adapted from Braverman et al.) Since $F \subseteq \{1,2,\ldots,T\}$, the proposition is equivalent to the following statement:

$$|F \cap \{1,2,\ldots,t\}| \leq \frac{t}{\Lambda}$$

for all $1 \leq t \leq T$. The proof proceeds by a potential argument. The potential function to consider is the sum of the credit values of Alice's edges. This quantity is always nonnegative and is initially zero. Any full page sent by Alice causes a decrement of the credit counter for each edge in the page, decreasing the potential function by Λ^2. On the other hand, any increase in the potential function is due to the arrival of a symbol (i.e., a good event or insertion addressed to Alice) and is precisely Λ. Since the potential function is nonnegative, we conclude that

$$\Lambda|F' \cap \{1,2,\ldots,t\}| \leq |(G' \cup I') \cap \{1,2,\ldots,t\}|.$$

Analogously,

$$\Lambda|F'' \cap \{1, 2, \ldots, t\}| \le |(G'' \cup I'') \cap \{1, 2, \ldots, t\}|.$$

Therefore,

$$\begin{aligned}\Lambda|F \cap \{1,2,\ldots,t\}|& \\ &\le \Lambda|F' \cap \{1, 2, \ldots, t\}| + \Lambda|F'' \cap \{1, 2, \ldots, t\}| \\ &\le |(G' \cup I') \cap \{1, 2, \ldots, t\}| + |(G'' \cup I'') \cap \{1, 2, \ldots, t\}|. \qquad (4.5.4)\end{aligned}$$

Parts (i), (ii), and (vii) of Lemma 4.21 imply that G', G'', I', I'' are pairwise disjoint. Therefore, the right-hand side of (4.5.4) does not exceed t. □

4.5.5. Finish times. We adopt the notation and definitions of Section 4.4.7, and review them here for the reader's convenience. Let $e_1, e_2, \ldots, e_n$ be the edges of the unique root-to-leaf path in $X \cup Y$, listed in increasing order of depth. Recall that a transmission *fulfills* an edge e if the corresponding page sent by the sender carries the last bit of an encoding of e. For $i = 1, 2, \ldots, n$, define f_i to be the index of the first transmission that fulfills e_i (whether or not that transmission is subject to an attack). If e_i is never fulfilled, we set $f_i = \infty$. For notational convenience, we also define $f_0 = f_{-1} = f_{-2} = \cdots = 0$. Recall from the description of the interactive coding scheme that Alice never starts sending an edge e unless she has finished sending all proper predecessors of e in X, and analogously for Bob. This gives

$$\begin{aligned}&f_1 \le f_3 \le f_5 \le \cdots, \\ &f_2 \le f_4 \le f_6 \le \cdots.\end{aligned}$$

The overall sequence $f_1, f_2, f_3, f_4, f_5, f_6, \ldots$ need not be in sorted order, however, due to interference by the adversary. We abbreviate

$$\overline{f_i} = \max\{0, f_1, f_2, \ldots, f_i\}.$$

By basic arithmetic,

$$[\overline{f_{i-1}}, \overline{f_i}) = [\overline{f_{i-1}}, f_i), \qquad\qquad i = 1, 2, \ldots, n. \tag{4.5.5}$$

Analogous to the analysis in Section 4.4.7, we need to bound the virtual length of each interval $[\overline{f_{i-1}}, f_i)$. To this end, we first bound the bit length of any encoding of e_i.

LEMMA 4.39. *For given integers i and t, suppose that an encoding of e_i is computed prior to the sending of transmission t. Then that encoding has bit length at most*

$$\lceil \log(t - f_{i-2}) \rceil + \lceil 2 \log \Lambda |\Sigma| \rceil.$$

Proof. Recall that e_i is encoded as a triple (m, j, σ), where m is the number of transmissions sent by the sender since his or her page that contained the last bit of e_{i-2} (for $i \geq 3$) or since the beginning of time (for $i = 1, 2$); $j \in \{1, 2, \ldots, \Lambda^2\}$ identifies e_{i-2} inside that page; and $\sigma \in \Sigma \times \Sigma$ identifies e_i relative to e_{i-2}. The pair (j, σ) takes on $\Lambda^2 |\Sigma|^2$ distinct values and can therefore be represented by a binary string of fixed length $\lceil 2 \log \Lambda |\Sigma| \rceil$. The remaining component m is a nonnegative integer of magnitude at most $t - f_{i-2} - 1$ and can therefore be represented by a binary string of length $\lceil \log(t - f_{i-2}) \rceil$ in the usual manner: $\varepsilon, 1, 10, 11, 100, \ldots$ for $0, 1, 2, 3, 4, \ldots$, respectively. □

We are now in a position to analyze the virtual length of any interval $[\overline{f_{i-1}}, f_i)$. The lemma that we are about to prove is a counterpart of Lemma 4.31.

Lemma 4.40. *For any $t \in \{1, 2, \ldots, T\}$ and i with $\overline{f_{i-1}} \le t < f_i$,*

$$\tilde{\|}[\overline{f_{i-1}}, t]\tilde{\|} \le \frac{2\Lambda}{\Lambda - 1}|B|_{[\overline{f_{i-1}},t]} + 2|F|_{[\overline{f_{i-1}},t]} \\ + 2\lceil \log(t - f_{i-2})\rceil + 2\lceil 2\log 2\Lambda|\Sigma|\rceil.$$

Proof. We will only treat the case of i odd; the proof for even i can be obtained by swapping the roles of Alice and Bob below.

For an edge e of the protocol tree, let *credit*(e, j) denote the value of *credit*(e) on Alice's side at the moment when transmission j enters the communication channel, i.e., immediately after the sender of transmission j has executed NextPage. For notational convenience, we also define *credit*$(e, 0) = 0$ for all e. Let $s \in [\overline{f_{i-1}}, t+1]$ be the smallest integer such that *credit*$(e_i, j) > 0$ for $j = s, s+1, \ldots, t$. Then

$$[\overline{f_{i-1}}, t] \subseteq [\overline{f_{i-1}}, s-1) \cup [s-1, s) \cup [s, t].$$

With this in mind, we complete the proof of the lemma by bounding the virtual length of each interval on the right-hand side and summing the resulting bounds. Key to our analysis are the following two claims.

Claim 4.41. $|E'|_{[\overline{f_{i-1}},s-1)} \le |B'|_{[\overline{f_{i-1}},s-1)}/(\Lambda - 1)$.

Proof. Consider any transmission $j \in E' \cap [\overline{f_{i-1}}, t)$. Lemma 4.37 ensures that on receipt of transmission j, Alice is able to correctly recover the complete set of edges that Bob has finished sending her by that time, which includes $e_2, e_4, e_6, \ldots, e_{i-1}$. At that same time, Alice has finished sending Bob $e_1, e_3, e_5, \ldots, e_{i-2}$ but not e_i, as one can verify from $\overline{f_{i-1}} \le j < t < f_i$. Therefore, the arrival of transmission j causes Alice to increase the credit of e_i by Λ in the call to AddEdge. The subsequent call to NextPage either leaves e_i's credit unchanged or decreases it by 1.

We now return to the proof of the claim. If $[\overline{f_{i-1}}, s-1) = \varnothing$, the claim holds trivially. In the complementary case, the definition of s ensures that

$$\mathit{credit}(e_i, s-1) = 0. \tag{4.5.6}$$

By the previous paragraph, the net effect of an incoming excellent transmission in the interval $[\overline{f_{i-1}}, t)$ is to increase e_i's credit by at least $\Lambda - 1$, whereas none of the other incoming symbols decrease e_i's credit by more than 1. Since credit is always nonnegative, we conclude from (4.5.6) that the number of incoming excellent transmissions in the interval $[\overline{f_{i-1}}, s-1)$ is at most a $1/(\Lambda - 1)$ fraction of the number of Alice's other incoming symbols in that interval. Formulaically, this conclusion translates to

$$|E'|_{[\overline{f_{i-1}}, s-1)} \leq \frac{1}{\Lambda - 1} |(G' \cup I') \setminus E'|_{[\overline{f_{i-1}}, s-1)},$$

which is equivalent to the claimed inequality by the definition of B'. □

Claim 4.42. $|G'' \cup D''|_{[s,t]} \leq |F'|_{[s,t]} + \lceil \log(t - f_{i-2}) \rceil + \lceil 2 \log \Lambda |\Sigma| \rceil$.

Proof. By the choice of s, the credit of e_i is positive when transmissions $s, s+1, \ldots, t$ enter the communication channel. Since Alice does not fulfill e_i before or during transmission $t < f_i$, we conclude that e_i is *continuously* present on Alice's edge list as transmissions $s, s+1, \ldots, t$ are prepared by their respective senders. In particular, every transmission among $s, s+1, \ldots, t$ that is sent by Alice must contain a bit of the encoding of e_i unless it is a full page. We conclude that

$$|G'' \cup D''|_{[s,t]} \leq |F'|_{[s,t]} + L,$$

where by definition $G'' \cup D''$ is the set of transmissions sent by Alice, F' is the set of transmissions sent by Alice that are full pages, and L stands for the bit length of e_i's encoding. This completes the proof in view of the upper bound on L in Lemma 4.39. □

It remains to put everything together. We have

$$\begin{aligned}
\tilde{\|}[\overline{f_{i-1}}, s-1)\tilde{\|} &\leq 2(|B|_{[\overline{f_{i-1}},s-1)} + |E'|_{[\overline{f_{i-1}},s-1)}) + 1 \\
&\leq 2\left(|B|_{[\overline{f_{i-1}},s-1)} + \frac{1}{\Lambda-1}|B'|_{[\overline{f_{i-1}},s-1)}\right) + 1 \\
&\leq \frac{2\Lambda}{\Lambda-1}|B|_{[\overline{f_{i-1}},s-1)} + 1,
\end{aligned} \tag{4.5.7}$$

where the first and second steps follow from Lemma 4.29 and Claim 4.41, respectively. Similarly,

$$\begin{aligned}
\tilde{\|}[s,t]\tilde{\|} &\leq 2(|B|_{[s,t]} + |E''|_{[s,t]}) + 1 \\
&\leq 2(|B|_{[s,t]} + |G''|_{[s,t]}) + 1 \\
&\leq 2(|B|_{[s,t]} + |F'|_{[s,t]} + \lceil \log(t - f_{i-2}) \rceil + \lceil 2 \log \Lambda |\Sigma| \rceil) + 1 \\
&\leq 2(|B|_{[s,t]} + |F|_{[\overline{f_{i-1}},t]} + \lceil \log(t - f_{i-2}) \rceil + \lceil 2 \log \Lambda |\Sigma| \rceil) + 1,
\end{aligned} \tag{4.5.8}$$

where the first and second steps follow from Lemma 4.29 and Claim 4.42, respectively. Finally,

$$\begin{aligned}
\tilde{\|}[s-1, s)\tilde{\|} &= \tilde{\|}\{s-1\}\tilde{\|} \\
&\leq 2.
\end{aligned} \tag{4.5.9}$$

Adding the bounds in (4.5.7)–(4.5.9) proves the lemma. □

4.5.6. The progress lemma. We have reached the technical centerpiece of our analysis, which is the counterpart of Lemma 4.32 for a large alphabet. Analogous to that earlier lemma, the result that we are about to prove shows that any sufficiently long execution of the interactive coding scheme with a sufficiently low virtual corruption rate allows Alice and Bob to exchange all the n edges of the unique root-to-leaf path in $X \cup Y$, and moreover this progress is not "undone" by any subsequent attacks by the adversary. Our exposition below emphasizes the similarities between Lemma 4.32 and the new result.

Lemma 4.43 (Progress lemma). *Let $t \in \{1, 2, \ldots, T\}$ be given with*

$$\tilde{\|}[1,t]\tilde{\|} \geq \frac{cn}{\alpha} \log \frac{|\Sigma|}{\alpha}, \tag{4.5.10}$$

$$\operatorname{corr}[1,t] \leq \frac{1}{4} - \alpha, \tag{4.5.11}$$

where $c \geq 1$ is a sufficiently large absolute constant. Then there is an integer $t^* \leq t$ such that

$$[\overline{f_n}, t^*) \cap E' \neq \varnothing, \tag{4.5.12}$$

$$[\overline{f_n}, t^*) \cap E'' \neq \varnothing, \tag{4.5.13}$$

$$\Delta(S'[i,t], R'[i,t]) < 1 - \alpha, \qquad i = 1, 2, \ldots, t^*, \tag{4.5.14}$$

$$\Delta(S''[i,t], R''[i,t]) < 1 - \alpha, \qquad i = 1, 2, \ldots, t^*. \tag{4.5.15}$$

Proof. Equations (4.5.14) and (4.5.15) hold vacuously for $t^* = 0$. In what follows, we will take $t^* \in \{0, 1, 2, \ldots, t\}$ to be the *largest* integer for which (4.5.14) and (4.5.15) hold. For the sake of contradiction, assume that at least one of the remaining desiderata (4.5.12), (4.5.13) is violated, whence

$$\tilde{\|}[\overline{f_n}, t^*)\tilde{\|} \leq 2|B|_{[\overline{f_n}, t^*)} + 1 \tag{4.5.16}$$

by Lemma 4.29. The proof strategy is to show that (4.5.16) is inconsistent with the hypothesis of the lemma. To this end, let $n^* \in \{0, 1, 2, \ldots, n\}$ be the largest integer such that $\overline{f_{n^*}} \leq t^*$. Then we have the partition

$$[0, t] = [\overline{f_0}, \overline{f_1}) \cup [\overline{f_1}, \overline{f_2}) \cup \cdots \cup [\overline{f_{n^*-1}}, \overline{f_{n^*}}) \cup [\overline{f_{n^*}}, t^*) \cup \{t^*\} \cup (t^*, t].$$

The bulk of our proof is concerned with bounding the virtual length of each of the intervals on the right-hand side.

Abbreviate

$$M = 2\lceil 2 \log 2\Lambda|\Sigma|\rceil + 2. \tag{4.5.17}$$

Then

$$\begin{aligned}\tilde{\|}[\overline{f_{i-1}}, \overline{f_i})\tilde{\|} &= \tilde{\|}[\overline{f_{i-1}}, f_i)\tilde{\|} \\ &\leq \frac{2\Lambda}{\Lambda - 1}|B|_{[\overline{f_{i-1}}, f_i)} + 2|F|_{[\overline{f_{i-1}}, f_i)} + 2\log(f_i - f_{i-2}) + M \\ &\leq \frac{2\Lambda}{\Lambda - 1}|B|_{[\overline{f_{i-1}}, \overline{f_i})} + 2|F|_{[\overline{f_{i-1}}, \overline{f_i})} + 2\log(f_i - f_{i-2}) + M \end{aligned} \tag{4.5.18}$$

for any $i = 1, 2, \ldots, n^*$, where the first and third steps use (4.5.5), and the second step follows from Lemma 4.40. Next, the upper bound

$$\begin{aligned}\tilde{\|}[\overline{f_{n^*}}, t^*)\tilde{\|} \leq \frac{2\Lambda}{\Lambda - 1}|B|_{[\overline{f_{n^*}}, t^*)} + 2|F|_{[\overline{f_{n^*}}, t^*)} \\ + 2\log(t^* - f_{n^*-1}) + M\end{aligned} \tag{4.5.19}$$

follows from Lemma 4.40 if $n^* < n$ and from (4.5.16) if $n^* = n$. The virtual length of the singleton interval $\{t^*\}$ can be bounded from first principles:

$$\tilde{\|}\{t^*\}\tilde{\|} \leq 2. \tag{4.5.20}$$

Finally, recall from the definition of t^* that either $\max\{\Delta(S'[t^*+1, t], R'[t^*+1, t]),$ $\Delta(S''[t^*+1, t], R''[t^*+1, t])\} \geq 1 - \alpha$ or $t^* = t$, leading to

$$\tilde{\|}(t^*, t]\tilde{\|} \leq \frac{4}{1 - \alpha}|D|_{(t^*, t]} + 1 \tag{4.5.21}$$

by Lemma 4.28 in the former case and trivially in the latter.

It remains to put together the upper bounds in (4.5.18)–(4.5.21). We have

$$\begin{aligned}
\tilde{\|}[1,t^*)\tilde{\|} &\le \frac{2\Lambda}{\Lambda-1}|B|_{[0,t^*)} + 2|F|_{[0,t^*)} + (n^*+1)M \\
&\qquad\qquad + 2\sum_{i=1}^{n^*}\log(f_i - f_{i-2}) + 2\log(t^* - f_{n^*-1}) \\
&\le \frac{2\Lambda}{\Lambda-1}\cdot\frac{2}{1-\alpha}|D|_{[0,t^*)} + \frac{2t^*}{\Lambda} + (n^*+1)M \\
&\qquad\qquad + 2\sum_{i=1}^{n^*}\log(f_i - f_{i-2}) + 2\log(t^* - f_{n^*-1}) \\
&\le \frac{2\Lambda}{\Lambda-1}\cdot\frac{2}{1-\alpha}|D|_{[0,t^*)} + \frac{2t^*}{\Lambda} + (n^*+1)M \\
&\qquad\qquad + 2(n^*+1)\log\frac{\sum_{i=1}^{n^*}(f_i - f_{i-2}) + (t^* - f_{n^*-1})}{n^*+1} \\
&= \frac{2\Lambda}{\Lambda-1}\cdot\frac{2}{1-\alpha}|D|_{[0,t^*)} + \frac{2t^*}{\Lambda} + (n^*+1)M \\
&\qquad\qquad + 2(n^*+1)\log\frac{f_{n^*} + t^*}{n^*+1} \\
&\le \frac{2\Lambda}{\Lambda-1}\cdot\frac{2}{1-\alpha}|D|_{[0,t^*)} + \frac{2t^*}{\Lambda} + (n^*+1)M \\
&\qquad\qquad + 2(n^*+1)\log\frac{2t^*}{n^*+1} \\
&\le \frac{4}{(1-\alpha)^2}|D|_{[0,t^*)} + \alpha\tilde{\|}[1,t]\tilde{\|} + 2(n^*+1)\cdot\left\lceil 1 + 2\log 2|\Sigma|\left\lceil\frac{2}{\alpha}\right\rceil\right\rceil \\
&\qquad\qquad + 2(n^*+1)\log\frac{2\tilde{\|}[1,t]\tilde{\|}}{n^*+1} \\
&\le \frac{4}{(1-\alpha)^2}|D|_{[0,t^*)} + 2\alpha\tilde{\|}[1,t]\tilde{\|} - 3, \qquad (4.5.22)
\end{aligned}$$

where the first step follows from (4.5.18) and (4.5.19); the second step applies Lemmas 4.25 and 4.38; the third step is valid by the concavity of the logarithm function; the next-to-last step is immediate from (4.5.2), (4.5.17), and $t^* \le t \le \tilde{\|}[1,t]\tilde{\|}$; and the last step follows from (4.5.10) and $n^* \le n$. Adding (4.5.20)–(4.5.22), we obtain

$$\tilde{\|}[1,t]\tilde{\|} \le \frac{4}{(1-\alpha)^2}|D|_{[0,t]} + 2\alpha\tilde{\|}[1,t]\tilde{\|},$$

or equivalently

$$\operatorname{corr}[1,t] \ge \frac{(1-2\alpha)(1-\alpha)^2}{4}.$$

This conclusion is inconsistent with (4.5.11) since $0 < \alpha < 1$. We have reached the desired contradiction and thereby proved that t^* satisfies each of the properties (4.5.12)–(4.5.15). □

4.5.7. Finishing the proof. We have reached a "master theorem" analogous to Theorem 4.33 for a large alphabet, which gives a sufficient condition for Alice and Bob to assign the correct value to their corresponding copies of the *out* variable. Once established, this new result will allow us to easily finish the proof of Theorem 4.34.

Theorem 4.44. *Consider a point in time when Alice updates her out variable, and fix a corresponding integer $t \le T$ such that $\not*(R'[1,t])$ is the complete sequence of symbols that Alice has received by that time. Assume that*

$$\tilde{\|}[1,t]\tilde{\|} \ge \frac{cn}{\alpha}\log\frac{|\Sigma|}{\alpha}, \tag{4.5.23}$$

$$\operatorname{corr}[1,t] \le \frac{1}{4} - \alpha, \tag{4.5.24}$$

where $c \ge 1$ is the absolute constant from Lemma 4.43*. Then as a result of the update, out is set to the leaf vertex in the unique root-to-leaf path in $X \cup Y$. An analogous theorem holds for Bob.*

Proof. Analogous to the proof of Theorem 4.33 for a large alphabet, with the difference that the newly obtained Lemmas 4.36 and 4.43 should be used instead of their large-alphabet counterparts (Lemmas 4.22 and 4.32). □

We now establish the main result of this chapter.

Proof of Theorem 4.34. The proof is nearly identical to that for a large alphabet (Theorem 4.19). Recall that $n = |\pi|$ denotes the communication cost of the original protocol, and $\varepsilon > 0$ is a constant in the statement of Theorem 4.19. Consider the interactive coding scheme given by Algorithms 5 and 6 with parameters set according to

$$\alpha = \frac{\varepsilon}{4}, \tag{4.5.25}$$

$$N = \left\lceil \frac{cn}{2\alpha} \log \frac{|\Sigma|}{\alpha} \right\rceil + 1, \tag{4.5.26}$$

where $c \geq 1$ is the absolute constant from Lemma 4.43. Then by (4.5.3), the interactive coding scheme uses an alphabet of size at most $2^{O(1/\varepsilon^3)} = O(1)$. Furthermore, the combined number of transmissions sent by Alice and Bob does not exceed $2N = O(\frac{n}{\varepsilon} \log \frac{|\Sigma|}{\varepsilon}) = O(|\pi| \log |\Sigma|)$.

It remains to show that when the communication stops, *out* is set for both Alice and Bob to the leaf vertex on the unique root-to-leaf path in $X \cup Y$. To this end, recall from (4.5.25) and Lemma 4.30 that

$$\operatorname{corr}[1, T] \leq \frac{1}{4} - 2\alpha. \tag{4.5.27}$$

By (4.5.26) and Lemma 4.27,

$$\tilde{\|}[1, T]\tilde{\|} \geq \frac{cn}{\alpha} \log \frac{|\Sigma|}{\alpha} + 2 \tag{4.5.28}$$

and therefore

$$\tilde{\|}[1, T-1]\tilde{\|} \geq \frac{cn}{\alpha} \log \frac{|\Sigma|}{\alpha}. \tag{4.5.29}$$

Also,

$$\begin{aligned}\mathrm{corr}[1,T-1] &\le \frac{\tilde{\|}[1,T]\tilde{\|}}{\tilde{\|}[1,T-1]\tilde{\|}}\cdot\mathrm{corr}[1,T]\\ &\le \left(1+\frac{2}{\tilde{\|}[1,T-1]\tilde{\|}}\right)\cdot\mathrm{corr}[1,T]\\ &\le \left(1+\frac{2\alpha}{n}\right)\cdot\left(\frac{1}{4}-2\alpha\right)\\ &\le \frac{1}{4}-\alpha, \end{aligned}\tag{4.5.30}$$

where the third step uses (4.5.27) and (4.5.29). Now, consider the last time that Alice and Bob update their copies of *out*. The complete sequence of symbols that Alice has received at the time of her last update is $\not\ast(R'[1,T-1])$ or $\not\ast(R'[1,T])$. Likewise, the complete sequence of symbols that Bob has received at the time of his last update is $\not\ast(R''[1,T-1])$ or $\not\ast(R''[1,T])$. By (4.5.27)–(4.5.30) and Theorem 4.44, both players set *out* to the leaf vertex in the unique root-to-leaf path in $X\cup Y$. □

4.5.8. Generalization to early output. Following Braverman et al. **[25]**, we now consider the setting when Alice and Bob need to be ready with their answers by a certain round (based on each player's own counting) rather than when the communication stops. Let Π be an interactive coding scheme. We define the δ*-early output* for a player in Π as the chronologically ordered sequence of the player's first $\delta|\Pi|/2$ symbols sent (or all of them, if the player sends fewer than $\delta|\Pi|/2$ symbols) and first $\delta|\Pi|/2$ symbols received (or all of them, if the player receives fewer than $\delta|\Pi|/2$ symbols). In this early output model, Alice and Bob are still expected to run their protocol to completion, which happens when one or both of them have finished $|\Pi|/2$ rounds of communication. The only difference is what information they use when computing their answers. Both Definition 4.10 and Theorem 4.34(i) generalize to the setting of early output, as follows.

Definition 4.45 (Coding scheme with early output). Let π be a given protocol with input space $\mathcal{X} \times \mathcal{Y}$. Protocol Π is an interactive coding scheme *for π with corruption rate ε and δ-early output* if:

(i) Π has input space $\mathcal{X} \times \mathcal{Y}$ and is in canonical form;

(ii) there are functions f', f'' such that for any pair of inputs $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ and any actions by an adversary with corruption rate ε, Alice's δ-early output a and Bob's δ-early output b satisfy $f'(a) = f''(b) = \pi(X, Y)$.

Theorem 4.46. *Fix arbitrary constants $\varepsilon > 0$ and $0 < \delta \le 1$, and let π be an arbitrary protocol with alphabet Σ. Then there exists an interactive coding scheme for π with alphabet size $O(1)$ and communication cost $O(|\pi| \log |\Sigma|)$ with δ-early output that tolerates corruption rate $(\frac{1}{4} - \varepsilon)\delta$.*

Proof. Let $n = |\pi|$ denote the communication cost of the original protocol. Consider the interactive coding scheme given by Algorithms 5 and 6 with parameters set according to

$$\alpha = \frac{\varepsilon}{2}, \tag{4.5.31}$$

$$N = \left\lceil \frac{cn}{\alpha\delta} \log \frac{|\Sigma|}{\alpha} + \frac{3}{\alpha\delta} \right\rceil, \tag{4.5.32}$$

where $c \ge 1$ is the absolute constant from Lemma 4.43. Then by (4.5.3), the interactive coding scheme uses an alphabet of size at most $2^{O(1/\varepsilon^3)} = O(1)$. Furthermore, the combined number of transmissions sent by Alice and Bob does not exceed $2N = O(\frac{n}{\varepsilon\delta} \log \frac{|\Sigma|}{\varepsilon}) = O(|\pi| \log |\Sigma|)$.

It remains to show that each player's δ-early output uniquely determines the output of π. We will prove the following much stronger statement: at any point in time when *one* of the players has processed $\delta|\Pi|/2 = \delta N$ or more incoming symbols, the variable *out* is set for both Alice and Bob to the leaf vertex of the unique root-to-leaf path in $X \cup Y$. This will prove the theorem since one of the players is always guaranteed to be able to run the protocol to completion and in particular to receive $|\Pi|/2 = N$ symbols.

We now provide the details. Fix any integer $t \in \{1, 2, \ldots, T\}$ such that at least one of the players receives $\delta|\Pi|/2$ or more symbols over the course of transmissions $1, 2, \ldots, t$. This is equivalent to saying that $\max\{|G' \cup I'|_{[1,t]}, |G'' \cup I''|_{[1,t]}\} \geq \delta N$. Lemma 4.26 implies that

$$\tilde{\|}[1,t]\tilde{\|} \geq 2\delta N - 1. \tag{4.5.33}$$

Now

$$\begin{aligned} \max\{\operatorname{corr}[1,t-1], \operatorname{corr}[1,t]\} &\leq \frac{|D|}{\tilde{\|}[1,t-1]\tilde{\|}} \\ &\leq \frac{\left(\frac{1}{4} - \varepsilon\right)\delta \cdot 2N}{2\delta N - 3} \\ &\leq \frac{1}{4} - \alpha, \end{aligned} \tag{4.5.34}$$

where the second step follows from the bound $|D| \leq (\frac{1}{4} - \varepsilon)\delta \cdot 2N$ in the hypothesis of the theorem, and the third step uses (4.5.31)–(4.5.33). Moreover, (4.5.32) and (4.5.33) ensure that

$$\tilde{\|}[1,t]\tilde{\|} \geq \tilde{\|}[1,t-1]\tilde{\|} \geq \frac{cn}{\alpha} \log \frac{|\Sigma|}{\alpha}. \tag{4.5.35}$$

Now, consider the last time that Alice and Bob update their copies of *out* over the course of transmissions $1, 2, \ldots, t$. The complete sequence of symbols that Alice has received at the time of her update is $\nparallel(R'[1,t-1])$ or $\nparallel(R'[1,t])$. Likewise, the complete sequence of symbols that Bob has received at the time of his update is $\nparallel(R''[1,t-1])$ or $\nparallel(R''[1,t])$. By (4.5.34), (4.5.35), and Theorem 4.44, both players set *out* to the leaf vertex in the unique root-to-leaf path in $X \cup Y$. □

In the terminology of our work, Braverman et al. [**25**] studied interactive coding schemes with $(1-2\eta)$-early output that tolerate corruption rate η. As their main result, they proved the existence of such a scheme with alphabet size $O(1)$ and communication cost $O(|\pi|\log|\Sigma|)$ for any constant $\eta < 1/18$. They also showed that no such scheme exists in general for $\eta \geq 1/6$. Our work closes the gap between the 1/18 and 1/6, establishing the existence of an interactive coding scheme for any π and any constant $\eta < 1/6$. This can be seen by taking $\delta = 1-2\eta$ and $\varepsilon = \frac{1}{4} - \frac{\eta}{1-2\eta}$ in Theorem 4.46.

4.5.9. Optimality. We now establish the optimality of Theorem 4.34, showing that it tolerates the highest possible corruption rate and normalized corruption rates. We do so by studying the *pointer jumping protocol* PJP_n, defined for $n \geq 1$ as the protocol with input space $\{0,1\}^n \times \{0,1\}^n$ in which Alice and Bob exchange their strings one bit at a time, taking turns after every bit. Thus, the sequence of symbols exchanged on input (x, y) is $x_1y_1 \ldots x_ny_n$. We show that no interactive coding scheme with alphabet size $2^{o(n)}$ for PJP_n can tolerate a corruption rate, or normalized corruption rates, higher than those tolerated by Theorem 4.34 with a constant-size alphabet. Our proof uses the "cut and paste" technique of previous impossibility results [**26, 25**]. We will first establish a detailed technical theorem and then deduce our impossibility results as corollaries.

THEOREM 4.47. *Let $\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}} \geq 0$ be given. Suppose that Π is an interactive coding scheme with alphabet Σ for PJP_n that tolerates normalized corruption rate $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$. Then*

$$\varepsilon_{\text{subs}} + \frac{3}{4}\varepsilon_{\text{oos}} < \frac{1}{4} + \frac{\log|\Sigma|}{n}. \tag{4.5.36}$$

Proof. Let $N = |\Pi|/2$ be the number of communication rounds in Π. Since Π simulates PJP_n, the former produces at least as many distinct transcripts as the latter. This leads to $|\Sigma|^{2N} \geq 4^n$ and

$$N \geq \frac{n}{\log |\Sigma|}. \tag{4.5.37}$$

The centerpiece of the proof is the following claim.

Claim 4.48. *The system*

$$k \leq \varepsilon_{\mathrm{oos}}(2N - k), \tag{4.5.38}$$

$$\left\lceil \frac{N}{2} \right\rceil - k \leq \varepsilon_{\mathrm{subs}}(2N - k) \tag{4.5.39}$$

has no integral solution $0 \leq k \leq \lceil N/2 \rceil$.

Proof. For the sake of contradiction, suppose that the system has a solution $k \in \{0, 1, 2, \ldots, \lceil N/2 \rceil\}$. Fix arbitrary $x, y, y' \in \{0,1\}^n$ with $y \neq y'$, and consider the following two executions of Π.

(i) Alice and Bob receive inputs x and y, respectively. The adversary uses substitution attacks to replace Bob's first $\lceil N/2 \rceil - k$ responses to Alice with the corresponding responses that he would send if his input were y'. Then the adversary carries out k consecutive out-of-sync attacks, intercepting Alice's transmissions to Bob and sending back to Alice the responses that Bob would send at that point if his input were y'. From then on, the adversary does not interfere with the communication. We let $\sigma_1, \sigma_2, \ldots, \sigma_N \in \Sigma$ denote the complete sequence of symbols that Alice receives in this execution.

(ii) Alice and Bob receive inputs x and y', respectively. The adversary does not interfere with the first $\lfloor N/2 \rfloor$ rounds of communication. As a result, the sequence of symbols that Alice receives in those rounds is $\sigma_1, \sigma_2, \ldots, \sigma_{\lfloor N/2 \rfloor}$. The adversary tampers with every symbol delivered to Alice from then on, making sure that she receives the sequence $\sigma_{\lfloor N/2 \rfloor + 1}, \ldots, \sigma_{N-1}, \sigma_N$. The adversary does so using $\lceil N/2 \rceil - k$ consecutive substitution attacks followed by k consecutive out-of-sync attacks. At that point, the communication stops because Alice has received N symbols.

Both executions feature $2N - k$ transmissions, $\lceil N/2 \rceil - k$ substitution attacks, and k out-of-sync attacks. By (4.5.38) and (4.5.39), these numbers of substitution and out-of-sync attacks are legitimate under normalized corruption rate $(\varepsilon_{\text{subs}}, \varepsilon_{\text{oos}})$. As a result, Alice and Bob's simulation of PJP_n is correct in both executions. Since PJP_n produces different transcripts on (x, y) and (x, y'), we conclude that both Alice and Bob are able to distinguish between the two executions. We have reached the promised contradiction because the two executions look identical to Alice. □

We now return to the proof of the theorem. The values $k \in [0, \lceil N/2 \rceil]$ that satisfy (4.5.38) form a subinterval of $[0, \lceil N/2 \rceil]$ that contains 0. Analogously, the values $k \in [0, \lceil N/2 \rceil]$ that satisfy (4.5.39) form a subinterval of $[0, \lceil N/2 \rceil]$ that contains $\lceil N/2 \rceil$. Since the system of these two inequalities has no integral solution in $[0, \lceil N/2 \rceil]$, there exists $k^* \in [0, \lceil N/2 \rceil - 1]$ such that $k = k^* + 1$ and $k = k^*$ violate (4.5.38) and (4.5.39), respectively:

$$\begin{aligned}
\varepsilon_{\text{oos}} &< \frac{k^* + 1}{2N - k^* - 1}, \\
\varepsilon_{\text{subs}} &< \frac{\lceil N/2 \rceil - k^*}{2N - k^*}.
\end{aligned}$$

Taking a weighted sum of these inequalities with weights 3/4 and 1,

$$\begin{aligned}
\frac{3}{4}\varepsilon_{\text{oos}} + \varepsilon_{\text{subs}} &< \frac{3}{4} \cdot \frac{k^* + 1}{2N - k^* - 1} + \frac{\frac{1}{2}(N+1) - k^*}{2N - k^*} \\
&= \frac{1}{4} + \frac{5N - k^* - 1}{2(2N - k^* - 1)(2N - k^*)} \\
&\leq \frac{1}{4} + \frac{1}{N},
\end{aligned}$$

where the last step uses $k^* \leq (N - 1)/2$. By (4.5.37), the proof is complete. □

We now derive our claimed impossibility results as corollaries of Theorem 4.47.

COROLLARY 4.49. *Suppose that for every $n \geq 1$, there is an interactive coding scheme for the pointer jumping protocol PJP_n with alphabet size $2^{o(n)}$ that tolerates normalized corruption rate $(\varepsilon_{\mathrm{subs}}, \varepsilon_{\mathrm{oos}})$. Then*

$$\varepsilon_{\mathrm{subs}} + \frac{3}{4}\varepsilon_{\mathrm{oos}} \leq \frac{1}{4}.$$

Proof. Substitute $|\Sigma| = 2^{o(n)}$ in Theorem 4.47 and pass to the limit as $n \to \infty$. □

COROLLARY 4.50. *Suppose that for every $n \geq 1$, there is an interactive coding scheme for the pointer jumping protocol PJP_n with alphabet size $2^{o(n)}$ that tolerates corruption rate ε. Then*

$$\varepsilon \leq \frac{1}{4}.$$

Proof. Any scheme that tolerates corruption rate ε must also tolerate normalized corruption rate $(\varepsilon, 0)$. Therefore, the claim follows by taking $\varepsilon_{\mathrm{subs}} = \varepsilon$ and $\varepsilon_{\mathrm{oos}} = 0$ in Corollary 4.49. □

CHAPTER 5

Unbounded-error communication complexity of $\mathbf{AC}^0$

In this chapter, we discuss our lower bound results in the strongest communication model, the unbounded-error communication complexity. Our proof relies on the analytic characterization described in Section 3.2.2. We first analyze the threshold degree of $\mathbf{AC}^0$. This implies lower bounds on communication with weakly unbounded error. Then we strengthen our analysis to that of sign rank of $\mathbf{AC}^0$, which in turn implies stronger lower bounds on communication with unbounded error.

5.1. Introduction

A real polynomial p is said to *sign-represent* the Boolean function $f\colon \{0,1\}^n \to \{0,1\}$ if $\operatorname{sgn} p(x) = (-1)^{f(x)}$ for every input $x \in \{0,1\}^n$. The *threshold degree* of f, denoted $\deg_{\pm}(f)$, is the minimum degree of a multivariate real polynomial that sign-represents f. Equivalent terms in the literature include *strong degree* [**10**], *voting polynomial degree* [**81**], *PTF degree* [**96**], and *sign degree* [**32**]. Since any function $f\colon \{0,1\}^n \to \{0,1\}$ can be represented exactly by a real polynomial of degree at most n, the threshold degree of f is an integer between 0 and n. Viewed as a computational model, sign-representation is remarkably powerful because it corresponds to the strongest form of pointwise approximation. The formal study of threshold degree began in 1969 with the pioneering work of Minsky and Papert [**90**] on limitations of perceptrons. The authors of [**90**] famously proved that the parity function on n variables has the maximum possible threshold degree,

n. They obtained lower bounds on the threshold degree of several other functions, including DNF formulas and intersections of halfspaces. Since then, sign-representing polynomials have found applications far beyond artificial intelligence. In theoretical computer science, applications of threshold degree include circuit lower bounds [**81, 82, 114, 41, 14**], size-depth trade-offs [**100, 132**], communication complexity [**114, 41, 116, 106, 14, 123, 122**], structural complexity theory [**16, 10**] and computational learning [**78, 77, 97, 9, 119, 121, 35, 122, 136**].

The notion of threshold degree has been especially influential in the study of $\mathbf{AC}^0$, the class of constant-depth polynomial-size circuits with $\wedge, \vee, \neg$ gates of unbounded fan-in. The first such result was obtained by Aspnes et al. [**10**], who used sign-representing polynomials to give a beautiful new proof of classic lower bounds for $\mathbf{AC}^0$. In communication complexity, the notion of threshold degree played a critical role in the first construction [**114, 116**] of an $\mathbf{AC}^0$ circuit with exponentially small discrepancy and hence large communication complexity in nearly every model. That discrepancy result was used in [**114**] to show the optimality of Allender's classic simulation of $\mathbf{AC}^0$ by majority circuits, solving the open problem [**81**] on the relation between the two circuit classes. Subsequent work [**42, 14, 125, 123**] resolved other questions in communication complexity and circuit complexity related to constant-depth circuits by generalizing the threshold degree method of [**114, 116**].

Sign-representing polynomials also paved the way for *algorithmic* breakthroughs in the study of constant-depth circuits. Specifically, any function of threshold degree d can be viewed as a halfspace in $\binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{d}$ dimensions, corresponding to the monomials in a sign-representation of f. As a result, a class of functions of threshold degree at most d can be learned in the standard PAC model under arbitrary distributions in time polynomial in $\binom{n}{0} + \binom{n}{1} + \cdots + \binom{n}{d}$. Klivans and Servedio [**78**] used this threshold degree approach to give what is currently the fastest algorithm for learning polynomial-size DNF formulas, with running time $\exp(\tilde{O}(n^{1/3}))$. Another learning-theoretic breakthrough based on threshold degree is the fastest algorithm for learning Boolean formulas, obtained by O'Donnell and Servedio [**97**] for formulas of constant depth and by Ambainis et al. [**9**] for arbitrary depth. Their algorithm runs in time $\exp(\tilde{O}(n^{(2^{k-1}-1)/(2^k-1)}))$ for formulas of size n and constant depth k, and in time $\exp(\tilde{O}(\sqrt{n}))$ for formulas of unbounded depth. In both cases, the bound on the running time follows from the corresponding upper bound on the threshold degree.

A far-reaching generalization of threshold degree is the matrix-analytic notion of *sign-rank*, which allows sign-representation out of arbitrary low-dimensional subspaces rather than the subspace of low-degree polynomials. The contribution of this chapter is to prove essentially optimal lower bounds on the threshold degree and sign-rank of $\mathbf{AC}^0$, which in turn imply lower bounds on other fundamental complexity measures of interest in communication complexity and learning theory. In the remainder of this section, we give a detailed overview of the previous work, present our main results, and discuss our proofs.

Depth	**Threshold degree**	**Reference**
2	$\Omega(n^{1/3})$	Minsky and Papert [**90**]
2	$O(n^{1/3}\log n)$	Klivans and Servedio [**78**]
k	$\Omega(n^{1/3}\log^{\frac{2(k-2)}{3}} n)$	O'Donnell and Servedio [**97**]
k	$\Omega(n^{\frac{k-1}{2k-1}})$	Sherstov [**122**]
4	$\Omega(\sqrt{n})$	Sherstov [**124**]
3	$\tilde{\Omega}(\sqrt{n})$	Bun and Thaler [**39**]
k	$\tilde{\Omega}(n^{\frac{k-1}{k+1}})$	This work

Table 1. Known bounds on the maximum threshold degree of $\wedge, \vee, \neg$-circuits of polynomial size and constant depth. In all bounds, n denotes the number of variables, and k denotes an arbitrary positive integer.

5.1.1. Threshold degree of $\mathbf{AC}^0$. Determining the maximum threshold degree of an $\mathbf{AC}^0$ circuit in n variables is a longstanding open problem in the area. It is motivated by the algorithmic and complexity-theoretic applications discussed above [**78, 97, 79, 106, 35**], in addition to being a natural question in its own right. Table 1 gives a quantitative summary of the results obtained to date. In their seminal monograph, Minsky and Papert [**90**] proved a lower bound of $\Omega(n^{1/3})$ on the threshold degree of the following DNF formula in n variables:

$$f(x) = \bigwedge_{i=1}^{n^{1/3}} \bigvee_{j=1}^{n^{2/3}} x_{i,j}.$$

Three decades later, Klivans and Servedio [**78**] obtained an $O(n^{1/3}\log n)$ upper bound on the threshold degree of any polynomial-size DNF formula in n variables, essentially matching Minsky and Papert's result and resolving the problem for depth 2. Determining the threshold degree of circuits of depth $k \geq 3$ proved to be challenging. The only upper bound known to date is the trivial $O(n)$, which follows directly from the definition of threshold degree. In particular, it is consistent with our knowledge that there are $\mathbf{AC}^0$ circuits with linear threshold degree. On the lower bounds side, the only progress for a long time was due to O'Donnell and Servedio [**97**], who constructed circuits of depth k with threshold degree $\Omega(n^{1/3}\log^{2(k-2)/3} n)$. The authors of [**97**] formally posed the problem of obtaining a polynomial improvement on Minsky and Papert's lower bound. Such an improvement was obtained in [**122**], with a threshold degree lower bound of $\Omega(n^{(k-1)/(2k-1)})$ for circuits of depth k. A polynomially stronger result was obtained in [**124**], with a lower bound of $\Omega(\sqrt{n})$ on the threshold degree of an explicit circuit of depth 4. Bun and Thaler [**39**] recently used a different, depth-3 circuit to give a much simpler proof of an $\tilde{\Omega}(\sqrt{n})$ lower bound for $\mathbf{AC}^0$. We obtain a quadratically stronger, and near-optimal, lower bound on the threshold degree of $\mathbf{AC}^0$.

Theorem 5.1. *Let $k \geq 1$ be a fixed integer. Then there is an (explicitly given) Boolean circuit family $\{f_n\}_{n=1}^{\infty}$, where $f_n \colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth k, and threshold degree*

$$\deg_{\pm}(f_n) = \Omega\left(n^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right).$$

Moreover, f_n has bottom fan-in $O(\log n)$ for all $k \neq 2$.

For large k, Theorem 5.1 essentially matches the trivial upper bound of n on the threshold degree of any function. For any fixed depth k, Theorem 5.1 subsumes all previous lower bounds on the threshold degree of $\mathbf{AC}^0$, with a polynomial improvement starting at depth $k = 4$. In particular, the lower bounds due to Minsky and Papert [**90**] and Bun and Thaler [**39**] are subsumed as the special cases $k = 2$ and $k = 3$, respectively. From a computational learning perspective, Theorem 5.1 definitively rules out the threshold degree approach to learning constant-depth circuits.

5.1.2. Sign-rank of $\mathbf{AC^0}$. The *sign-rank* of a matrix $A = [A_{ij}]$ without zero entries is the least rank of a real matrix $M = [M_{ij}]$ with $\operatorname{sgn} M_{ij} = \operatorname{sgn} A_{ij}$ for

Depth	**Sign-rank**	**Reference**
3	$\exp(\Omega(n^{1/3}))$	Razborov and Sherstov [**106**]
3	$\exp(\Omega(n^{2/5}))$	Bun and Thaler [**37**]
7	$\exp(\tilde{\Omega}(\sqrt{n}))$	Bun and Thaler [**39**]
$3k$	$\exp(\tilde{\Omega}(n^{1-\frac{1}{k+1}}))$	This work
$3k+1$	$\exp(\tilde{\Omega}(n^{1-\frac{1}{k+1.5}}))$	This work

Table 2. Known lower bounds on the maximum sign-rank of $\wedge, \vee, \neg$-circuits of polynomial size and constant depth. In all bounds, n denotes the number of variables, and k denotes an arbitrary positive integer.

all i, j. In other words, the sign-rank of A is the minimum rank of a matrix that can be obtained by making arbitrary sign-preserving changes to the entries of A. The sign-rank of a Boolean function $F\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ is defined in the natural way as the sign-rank of the matrix $[(-1)^{F(x,y)}]_{x,y}$. In particular, the sign-rank of F is an integer between 1 and 2^n. This fundamental notion has been studied in contexts as diverse as matrix analysis, communication complexity, circuit complexity, and learning theory [**101, 7, 17, 53, 54, 78, 88, 113, 117, 106, 37, 39**]. To a complexity theorist, sign-rank is a vastly more challenging quantity to analyze than threshold degree. Indeed, a sign-rank lower bound rules out a sign-representation out of *every* linear subspace of given dimension, whereas a threshold degree lower bound rules out a sign-representation specifically by linear combinations of monomials up to a given degree.

Unsurprisingly, progress in understanding sign-rank has been slow and difficult. No nontrivial lower bounds were available for any explicit matrices until the breakthrough work of Forster [**53**], who proved strong lower bounds on the sign-rank of Hadamard matrices and more generally all sign matrices with small spectral norm. The sign-rank of constant-depth circuits $F\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ has since seen considerable work, as summarized in Table 2. The first exponential lower bound on the sign-rank of an $\mathbf{AC}^0$ circuit was obtained by Razborov and Sherstov [**106**], solving a 22-year-old problem due to Babai, Frankl, and Simon [**11**]. The authors of [**106**] constructed a polynomial-size circuit of depth 3 with sign-rank $\exp(\Omega(n^{1/3}))$. In follow-up work, Bun and Thaler [**37**] constructed a polynomial-size circuit of depth 3 with sign-rank $\exp(\tilde{\Omega}(n^{2/5}))$. A more recent and incomparable result, also due to Bun and Thaler [**39**], is a sign-rank lower bound of $\exp(\tilde{\Omega}(\sqrt{n}))$ for a circuit of polynomial size and depth 7. No nontrivial upper bounds are known on the sign-rank of $\mathbf{AC}^0$. Closing this gap between the best lower bound of $\exp(\tilde{\Omega}(\sqrt{n}))$ and the trivial upper bound of 2^n has been a challenging open problem. We solve this problem almost completely, by constructing for any $\varepsilon > 0$ a constant-depth circuit with sign-rank $\exp(\Omega(n^{1-\varepsilon}))$. In quantitative detail, our results on the sign-rank of $\mathbf{AC}^0$ are the following two theorems.

THEOREM 5.2. *Let $k \geq 1$ be a given integer. Then there is an (explicitly given) Boolean circuit family $\{F_n\}_{n=1}^{\infty}$, where $F_n\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k$, and sign-rank*

$$\mathrm{rk}_{\pm}(F_n) = \exp\left(\Omega\left(n^{1-\frac{1}{k+1}} \cdot (\log n)^{-\frac{k(k-1)}{2(k+1)}}\right)\right).$$

As a companion result, we prove the following qualitatively similar but quantitatively incomparable theorem.

THEOREM 5.3. *Let $k \geq 1$ be a given integer. Then there is an (explicitly given) Boolean circuit family $\{G_n\}_{n=1}^{\infty}$, where $G_n\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k+1$, and sign-rank*

$$\mathrm{rk}_{\pm}(G_n) = \exp\left(\Omega\left(n^{1-\frac{1}{k+1.5}} \cdot (\log n)^{-\frac{k^2}{2k+3}}\right)\right).$$

For large k, the lower bounds of Theorems 5.2 and 5.3 approach the trivial upper bound of 2^n on the sign-rank of any Boolean function $\{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. For any fixed depth k, Theorems 5.2 and 5.3 subsume all previous lower bounds on the sign-rank of $\mathbf{AC}^0$, with a strict improvement starting at depth 3. From a computational learning perspective, Theorems 5.2 and 5.3 state that $\mathbf{AC}^0$ has near-maximum *dimension complexity* [**113, 115, 106, 39**], namely, $\exp(\Omega(n^{1-\varepsilon}))$ for any constant $\varepsilon > 0$. This rules out the possibility of learning $\mathbf{AC}^0$ circuits via dimension complexity [**106**], a far-reaching generalization of the threshold degree approach from the monomial basis to arbitrary bases.

5.1.3. Communication complexity. Theorems 5.1–5.3 imply strong new lower bounds on the communication complexity of $\mathbf{AC}^0$. To begin with, combining Theorem 5.1 with the *pattern matrix method* [**114, 116**] gives:

THEOREM 5.4. *Let $k \geq 3$ be a fixed integer. Then there is an (explicitly given) Boolean circuit family $\{F_n\}_{n=1}^{\infty}$, where $F_n\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ has polynomial size, depth k, communication complexity*

$$\mathsf{PP}(F_n) = \Omega\left(n^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right)$$

and discrepancy

$$\operatorname{disc}(F_n) = \exp\left(-\Omega\left(n^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right)\right).$$

Discrepancy is a combinatorial complexity measure of interest in communication complexity theory and other research areas; see Section 3.2.2 for a formal definition. As k grows, the bounds of Theorem 5.4 approach the best possible bounds for any communication problem $F\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. The same *qualitative* behavior was achieved in previous work by Bun and Thaler [**39**], who constructed, for any constant $\varepsilon > 0$, a constant-depth circuit $F_n\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ with communication complexity $\mathsf{PP}(F) = \Omega(n^{1-\varepsilon})$ and discrepancy $\operatorname{disc}(F) = \exp(-\Omega(n^{1-\varepsilon}))$. Theorem 5.4 strictly subsumes the result of Bun and Thaler [**39**] and all other prior work on the discrepancy and PP-complexity of constant-depth circuits [**114, 116, 14, 125, 123**]. For any fixed depth $k \geq 4$, the bounds of Theorem 5.4 are a polynomial improvement in n over all previous work. We further obtain a counterpart of Theorem 5.4 for *number-on-the-forehead model*, the strongest formalism of multiparty communication. This result, presented in detail in Section 5.4.5, uses the multiparty version [**123**] of the pattern matrix method.

Our work also gives near-optimal lower bounds for $\mathbf{AC}^0$ in the much more powerful unbounded-error model. Specifically, it is well-known [**101**] that the unbounded-error communication complexity of any Boolean function $F\colon X \times Y \to \{0,1\}$ coincides up to an additive constant with the logarithm of the sign-rank of F. As a result, Theorems 5.2 and 5.3 imply:

THEOREM 5.5. *Let* $k \geq 1$ *be a given integer. Let* $\{F_n\}_{n=1}^\infty$ *and* $\{G_n\}_{n=1}^\infty$ *be the polynomial-size circuit families of depth* $3k$ *and* $3k+1$, *respectively, constructed in Theorems* 5.2 *and* 5.3*. Then*

$$\begin{aligned}\mathsf{UPP}(F_n) &= \Omega\left(n^{1-\frac{1}{k+1}} \cdot (\log n)^{-\frac{k(k-1)}{2(k+1)}}\right),\\ \mathsf{UPP}(G_n) &= \Omega\left(n^{1-\frac{1}{k+1.5}} \cdot (\log n)^{-\frac{k^2}{2k+3}}\right).\end{aligned}$$

For large k, the lower bounds of Theorem 5.5 essentially match the trivial upper bound of $n+1$ on the unbounded-error communication complexity of any function $F\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$. Theorem 5.5 strictly subsumes all previous work on the unbounded-error communication complexity of $\mathbf{AC}^0$, with a polynomial improvement for any depth $k \geq 3$. The best lower bound on the unbounded-error communication complexity of $\mathbf{AC}^0$ prior to our work was $\tilde{\Omega}(\sqrt{n})$ for a circuit of depth 7, due to Bun and Thaler [**39**]. Finally, we remark that Theorem 5.5 gives essentially the strongest possible separation of the communication complexity classes PH and UPP. We refer the reader to the work of Babai et al. [**11**] for definitions and detailed background on these classes.

Qualitatively, Theorem 5.5 is stronger than Theorem 5.4 because communication protocols with unbounded error are significantly more powerful than those with weakly unbounded error. On the other hand, Theorem 5.4 is stronger quantitatively for any fixed depth k and has the additional advantage of generalizing to the multiparty setting.

5.1.4. Threshold weight and threshold density. By well-known reductions, Theorem 5.1 implies a number of other lower bounds for the representation of $\mathbf{AC}^0$ circuits by polynomials. For the sake of completeness, we mention two such consequences. The *threshold density* of a Boolean function $f\colon \{0,1\}^n \to \{0,1\}$, denoted $\operatorname{dns}(f)$, is the minimum size of a set family $\mathcal{S} \subseteq \mathcal{P}(\{1,2,\dots,n\})$ such that

$$\operatorname{sgn}\left(\sum_{S\in\mathcal{S}} \lambda_S (-1)^{\sum_{i\in S} x_i}\right) \equiv (-1)^{f(x)}$$

for some reals λ_S. A related complexity measure is *threshold weight*, denoted $W(f)$ and defined as the minimum sum $\sum_{S\subseteq\{1,2,\dots,n\}} |\lambda_S|$ over all integers λ_S such that

$$\operatorname{sgn}\left(\sum_{S\subseteq\{1,2,\dots,n\}} \lambda_S (-1)^{\sum_{i\in S} x_i}\right) \equiv (-1)^{f(x)}.$$

It is not hard to see that the threshold density and threshold weight of f correspond to the minimum size of a threshold-of-parity and majority-of-parity circuit for f, respectively. The definitions imply that $\operatorname{dns}(f) \leq W(f)$ for every f, and a little more thought reveals that $1 \leq \operatorname{dns}(f) \leq 2^n$ and $1 \leq W(f) \leq (2\sqrt{2})^n$. These complexity measures have seen extensive work, motivated by applications to computational learning and circuit complexity. For a bibliographic overview, we refer the reader to [**122**, Section 8.2].

Krause and Pudlák [**81**, Proposition 2.1] gave an ingenious method for transforming threshold degree lower bounds into lower bounds on threshold density and thus also threshold weight. Specifically, let $f\colon \{0,1\}^n \to \{0,1\}$ be a Boolean function of interest. The authors of [**81**] considered the related function $F\colon (\{0,1\}^n)^3 \to \{0,1\}$ given by $F(x,y,z) = f(\dots,(\overline{z_i} \wedge x_i) \vee (z_i \wedge y_i),\dots)$, and proved that $\operatorname{dns}(F) \geq 2^{\deg_\pm(f)}$. In this light, Theorem 5.1 implies that the threshold density of $\mathbf{AC}^0$ is $\exp(\Omega(n^{1-\varepsilon}))$ for any constant $\varepsilon > 0$.

COROLLARY 5.6. *Let $k \geq 3$ be a fixed integer. Then there is an (explicitly given) Boolean circuit family $\{F_n\}_{n=1}^{\infty}$, where $F_n \colon \{0,1\}^n \to \{0,1\}$ has polynomial size and depth k and satisfies*

$$\begin{aligned} W(F_n) &\geq \operatorname{dns}(F_n) \\ &= \exp\left(\Omega\left(n^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right)\right). \end{aligned}$$

Observe that the circuit family $\{F_n\}_{n=1}^{\infty}$ of Corollary 5.6 has the same depth as the circuit family $\{f_n\}_{n=1}^{\infty}$ of Theorem 5.1. This is because f_n has bottom fan-in $O(\log n)$, and thus the Krause-Pudlák transformation $f_n \mapsto F_n$ can be "absorbed" into the bottom two levels of f_n. Corollary 5.6 subsumes all previous lower bounds [**81, 35, 122, 124, 39**] on the threshold weight and density of $\mathbf{AC}^0$, with a polynomial improvement for every $k \geq 4$. The improvement is particularly noteworthy in the case of threshold density, where the best previous lower bound [**124, 39**] was $\exp(\Omega(\sqrt{n}))$.

5.1.5. Previous approaches. In the remainder of this section, we discuss our proofs of Theorems 5.1–5.3. The notation that we use here is standard, and we defer its formal review to Section 5.2. We start with necessary approximation-theoretic background, then review relevant previous work, and finally contrast it with the approach of this paper. To sidestep minor technicalities, we will represent Boolean functions in this overview as mappings $\{-1,1\}^n \to \{-1,1\}$. We alert the reader that we will revert to the standard $\{0,1\}^n \to \{0,1\}$ representation starting with Section 5.2.

Background. Recall that our results concern the sign-representation of Boolean functions and matrices. To properly set the stage for our proofs, however, we need to consider the more general notion of pointwise approximation [**93**]. Let $f\colon \{-1,1\}^n \to \{-1,1\}$ be a Boolean function of interest. The *ε-approximate degree of f*, denoted $\deg_\varepsilon(f)$, is the minimum degree of a real polynomial that approximates f within ε pointwise: $\deg_\varepsilon(f) = \min\{\deg p\colon \|f-p\|_\infty \leq \varepsilon\}$. The regimes of most interest are *bounded-error approximation*, corresponding to constants $\varepsilon \in (0,1)$; and *large-error approximation*, corresponding to $\varepsilon = 1-o(1)$. In the former case, the choice of the error parameter $\varepsilon \in (0,1)$ is immaterial and affects the approximate degree of a Boolean function by at most a multiplicative constant. It is clear that pointwise approximation is a stronger requirement than sign-representation, and thus $\deg_\pm(f) \leq \deg_\varepsilon(f)$ for all $0 \leq \varepsilon < 1$. A moment's thought reveals that threshold degree is in fact the limiting case of ε-approximate degree as the error parameter approaches 1:

$$\deg_\pm(f) = \lim_{\varepsilon \nearrow 1} \deg_\varepsilon(f). \tag{5.1.1}$$

Both approximate degree and threshold degree have dual characterizations [**116**], obtained by appeal to linear programming duality. Specifically, $\deg_\varepsilon(f) \geq d$ if and only if there is a function $\phi\colon \{-1,1\}^n \to \mathbb{R}$ with the following two properties: $\langle \phi, f\rangle > \varepsilon\|\phi\|_1$; and $\langle \phi, p\rangle = 0$ for every polynomial of degree less than d. Rephrasing, ϕ must have large correlation with f but zero correlation with every low-degree polynomial. By weak linear programming duality, ϕ constitutes a proof that $\deg_\varepsilon(f) \geq d$ and for that reason is said to *witness* the lower bound $\deg_\varepsilon(f) \geq d$. In view of (5.1.1), this discussion carries over to the case of threshold degree. The dual characterization here states that $\deg_\pm(f) \geq d$ if and only if there is a nonzero function $\phi\colon \{-1,1\}^n \to \mathbb{R}$ with the following two properties: $\phi(x)f(x) \geq 0$ for all x; and $\langle \phi, p\rangle = 0$ for every polynomial of degree less than d. In this dual characterization, ϕ agrees in sign with f and is additionally orthogonal to polynomials of degree less than d. The sign-agreement property can be restated in terms of correlation, as $\langle \phi, f \rangle = \|\phi\|_1$. As before, ϕ is called a threshold degree *witness* for f.

What distinguishes the dual characterizations of approximate degree and threshold degree is how the dual object ϕ relates to f. Specifically, a threshold degree witness must agree in sign with f at every point. An approximate degree witness, on the other hand, need only exhibit such sign-agreement with f at *most* points, in that the points where the sign of ϕ is correct should account for most of the ℓ_1 norm of ϕ. As a result, constructing dual objects for threshold degree is significantly more difficult than for approximate degree. This difficulty is to be expected because because the gap between threshold degree and approximate degree can be arbitrary, e.g., 1 versus $\Theta(n)$ for the majority function on n bits [**99**].

Hardness amplification via block-composition. Much of the recent work on approximate degree and threshold degree is concerned with composing functions in ways that amplify their hardness. Of particular significance here is *block-composition*, defined for functions $f\colon \{-1,1\}^n \to \{-1,1\}$ and $g\colon X \to \{-1,1\}$ as the Boolean function $f \circ g\colon X^n \to \{-1,1\}$ given by $(f \circ g)(x_1, \ldots, x_n) = f(g(x_1), \ldots, g(x_n))$. Block-composition works particularly well for threshold degree. To use an already familiar example, the block-composition $\mathrm{AND}_{n^{1/3}} \circ \mathrm{OR}_{n^{2/3}}$ has threshold degree $\Omega(n^{1/3})$ whereas the constituent functions $\mathrm{AND}_{n^{1/3}}$ and $\mathrm{OR}_{n^{2/3}}$ have threshold degree 1. As a more extreme example, Sherstov [**121**] obtained a lower bound of $\Omega(n)$ on the threshold degree of the conjunction $h_1 \wedge h_2$ of two halfspaces $h_1, h_2\colon \{0,1\}^n \to \{0,1\}$, each of which by definition has threshold degree 1. The fact that threshold degree can increase spectacularly under block-composition was the basis of much previous work, including the best previous lower bounds [**122, 124**] on the threshold degree of $\mathbf{AC}^0$. Apart from threshold degree, block-composition has yielded strong results for approximate degree in various error regimes, including direct sum theorems [**119**] and direct product theorems [**118**] for approximate degree and error amplification for approximate degree [**118, 35, 136, 36**].

How, then, does one prove lower bounds on the threshold degree or approximate degree of a composed function $f \circ g$? It is here that the dual characterizations take center stage: they make it possible to prove lower bounds *algorithmically*, by constructing the corresponding dual object ϕ for the function of interest. Such algorithmic proofs run the gamut in terms of technical sophistication, from straightforward to lengthy and highly technical, but they have some structure in common. In most cases, one starts by obtaining dual objects ϕ and ψ for the constituent functions f and g, respectively, either by direct construction or by appeal to linear programming duality. They are then combined to yield a dual object Φ for the composed function, using *dual block-composition* [**119, 85**]:

$$\Phi(x_1, x_2, \ldots, x_n) = \phi(\operatorname{sgn}\psi(x_1), \ldots, \operatorname{sgn}\psi(x_n)) \prod_{i=1}^{n} |\psi(x_i)|. \tag{5.1.2}$$

This composed dual object often requires additional work to ensure sign-agreement or correlation with the composed Boolean function. Among the generic tools available to assist in this process is a "corrector" object ζ due to Razborov and Sherstov [**106**], with the following four properties: (i) ζ is orthogonal to low-degree polynomials; (ii) ζ takes on 1 at a prescribed point of the hypercube; (iii) ζ is bounded on inputs of low Hamming weight; and (iv) ζ vanishes on all other points of the hypercube. Using the Razborov–Sherstov object, suitably shifted and scaled, one can surgically correct the behavior of a given dual object Φ on a substantial fraction of inputs, thus modifying its metric properties without affecting its orthogonality to low-degree polynomials. This technique has played an important role in recent work, e.g., [**37, 38, 33, 39**].

Hardness amplification for approximate degree. While block-composition has produced a treasure trove of results on the polynomial representation of Boolean functions, it is of limited use when it comes to constructing functions with high *bounded-error* approximate degree. To illustrate the issue, consider arbitrary functions $f\colon \{-1,1\}^{n_1} \to \{-1,1\}$ and $g\colon \{-1,1\}^{n_2} \to \{-1,1\}$ with 1/3-approximate degrees $n_1^{\alpha_1}$ and $n_2^{\alpha_2}$, respectively, for some $0 < \alpha_1 < 1$ and $0 < \alpha_2 < 1$. It is well-known [**120**] that the composed function $f \circ g$ on $n_1 n_2$ variables has 1/3-approximate degree $O(n_1^{\alpha_1} n_2^{\alpha_2}) = O(n_1 n_2)^{\max\{\alpha_1,\alpha_2\}}$. This means that relative to the new number of variables, the block-composed function $f \circ g$ is no harder to approximate to bounded error than either of the constituent functions f and g. In particular, one cannot use block-composition to transform functions on n bits with 1/3-approximate degree at most n^α into functions on $N \geq n$ bits with 1/3-approximate degree $\omega(N^\alpha)$.

Until recently, the best lower bound on the bounded-error approximate degree of $\mathbf{AC}^0$ was $\Omega(n^{2/3})$, due to Aaronson and Shi [**5**]. Breaking this $n^{2/3}$ barrier was a fundamental problem in its own right, in addition to being a hard prerequisite for any future *threshold* degree lower bounds for $\mathbf{AC}^0$ better than $\Omega(n^{2/3})$. This barrier was overcome in a brilliant paper of Bun and Thaler [**38**], who proved, for any constant $\varepsilon > 0$, an $\Omega(n^{1-\varepsilon})$ lower bound on the 1/3-approximate degree of $\mathbf{AC}^0$. In more detail, let $f\colon \{-1,1\}^n \to \{-1,1\}$ be a function of interest, with 1/3-approximate degree n^α for some $0 \leq \alpha < 1$. Bun and Thaler consider the block-composition $F = f \circ \mathrm{AND}_{\Theta(\log m)} \circ \mathrm{OR}_m$, for an appropriate parameter $m = \mathrm{poly}(n)$. As shown in earlier work [**119, 35**] on approximate degree, dual block-composition witnesses the lower bound $\deg_{1/3}(F) = \Omega(\deg_{1/3}(\mathrm{OR}_m)\deg_{1/3}(f)) = \Omega(\sqrt{m}\deg_{1/3}(f))$. Here, Bun and Thaler make the crucial observation that the dual object for OR_m has most of its ℓ_1 mass on inputs of Hamming weight $O(1)$, which in view of (5.1.2) implies that the dual object for F places most of it ℓ_1 mass on inputs of Hamming weight

$O(n \log n)$. The authors of [**38**] then use the Razborov–Sherstov corrector object to transfer the small amount of ℓ_1 mass that the dual object for F places on inputs of high Hamming weight, to inputs of low Hamming weight. The resulting dual object for F is supported entirely on inputs of low Hamming weight and therefore witnesses a lower bound on the 1/3-approximate degree of the *restriction* F' of F to inputs of low Hamming weight. By re-encoding the input to F', one finally obtains a function F'' on $n(\log n)^{O(1)}$ variables with 1/3-approximate degree polynomially larger than that of f. This passage from f to F'' is the desired hardness amplification for approximate degree. We find it helpful to think of Bun and Thaler's technique as block-composition followed by input compression, to reduce the number of input variables in the block-composed function. To obtain an $\Omega(n^{1-\varepsilon})$ lower bound on the approximate degree of $\mathbf{AC}^0$, the authors of [**38**] start with a trivial circuit and iteratively apply the hardness amplification step a constant number of times, until approximate degree $\Omega(n^{1-\varepsilon})$ is reached.

In follow-up work, Bun, Kothari, and Thaler [**33**] refined the technique of [**38**] by deriving optimal concentration bounds for the dual object for OR_m. They thereby obtained tight lower bounds on the 1/3-approximate degree of *surjectivity*, *element distinctness*, and other important problems. The most recent contribution to this line of work is due to Bun and Thaler [**39**], who prove an $\Omega(n^{1-\varepsilon})$ lower bound on the $(1 - 2^{-n^{1-\varepsilon}})$-approximate degree of $\mathbf{AC}^0$ by combining the method of [**38**] with Sherstov's work [**118**] on direct product theorems for approximate degree. This new result substantially strengthens the authors' previous result [**38**] on the *bounded-error* approximate degree of $\mathbf{AC}^0$ but falls short of a threshold degree lower bound.

5.1.6. Our approach.

Threshold degree of AC^0. Bun and Thaler [**39**] refer to obtaining an $\Omega(n^{1-\varepsilon})$ threshold degree lower bound for $\mathbf{AC}^0$ as the "main glaring open question left by our work." It is important to note here that lower bounds on approximate degree, even with the error parameter exponentially close to 1 as in [**39**], have no implications for threshold degree. For example, there are functions [**121**] with $(1-2^{-\Theta(n)})$-approximate degree $\Theta(n)$ but threshold degree 1. Our proof of Theorem 5.1 is unrelated to the most recent work of Bun and Thaler [**39**] on the large-error approximate degree of $\mathbf{AC}^0$ and instead builds on the earlier and simpler "block-composition followed by input compression" approach of [**38**]. The centerpiece of our proof is a hardness amplification result for threshold degree, whereby any function f with threshold degree n^α for a constant $0 \leq \alpha < 1$ is transformed efficiently and within $\mathbf{AC}^0$ into a function F with polynomially larger threshold degree.

In more detail, let $f\colon \{-1,1\}^n \to \{-1,1\}$ be a function of interest, with threshold degree n^α. We consider the block-composition $f \circ \mathrm{MP}_m$, where $m = n^{O(1)}$ is an appropriate parameter and $\mathrm{MP}_m = \mathrm{AND}_m \circ \mathrm{OR}_{m^2}$ is the Minsky–Papert function with threshold degree $\Omega(m)$. We construct the dual object for MP_m from scratch to ensure concentration on inputs of Hamming weight $\tilde{O}(m)$. By applying dual block-composition to the threshold degree witnesses of f and MP_m, we obtain a dual object Φ witnessing the $\Omega(mn^\alpha)$ threshold degree of $f \circ \mathrm{MP}_m$. So far in the proof, our differences from [**38**] are as follows: (i) since our goal is amplification of threshold degree, we work with witnesses of threshold degree rather than approximate degree; (ii) to ensure rapid growth of threshold degree, we use block-composition with inner function $\mathrm{MP}_m = \mathrm{AND}_m \circ \mathrm{OR}_{m^2}$ of threshold degree $\Theta(m)$, in place of Bun and Thaler's inner function $\mathrm{AND}_{\Theta(\log m)} \circ \mathrm{OR}_m$ of threshold degree $\Theta(\log m)$.

Since the dual object for MP_m by construction has most of its ℓ_1 norm on inputs of Hamming weight $\tilde{O}(m)$, the dual object Φ for the composed function has most of its

ℓ_1 norm on inputs of Hamming weight $\tilde{O}(nm)$. Analogous to [**38, 33, 39**], we would like to use the Razborov–Sherstov corrector object to *remove* the ℓ_1 mass that Φ has on inputs on high Hamming weight, transferring it to inputs of low Hamming weight. This brings us to the novel and technically demanding part of our proof. Previous works [**38, 33, 39**] transferred the ℓ_1 mass from inputs of high Hamming weight to the neighborhood of the all-zeroes input $(0, 0, \ldots, 0)$. An unavoidable downside of the Razborov–Sherstov transfer process is that it amplifies the ℓ_1 mass being transferred. When the transferred mass finally reaches its destination, it overwhelms Φ's original values at various points, destroying Φ's sign-agreement with the composed function $f \circ \mathrm{MP}_m$. It is this difficulty that prevented earlier works [**38, 33, 39**] from obtaining a strong threshold degree lower bound for $\mathbf{AC}^0$.

We proceed differently. Instead of transferring the ℓ_1 mass of Φ from inputs of high Hamming weight to the neighborhood of $(0, 0, \ldots, 0)$, we transfer it simultaneously to *exponentially many* neighborhoods of inputs with low Hamming weight. Split this way across many neighborhoods, the transferred mass does not overpower the original values of Φ and in particular does not change any signs. Working out the details of this transfer scheme requires subtle calculations; it is in fact surprising that such a scheme exists. Once the transfer process is complete, we obtain a witness for the $\Omega(mn^\alpha)$ threshold degree of $f \circ \mathrm{MP}_m$ even for the restriction of the domain to inputs of low Hamming weight. Compressing the input as in [**38, 33**], we obtain an amplification theorem for threshold degree. With this work behind us, the proof of Theorem 5.1 for any depth k amounts to starting with a trivial circuit and amplifying its threshold degree $O(k)$ times.

Sign-rank of AC^0. It is not known how to transform a threshold degree lower bound in a black-box manner into a sign-rank lower bound. In particular, Theorem 5.1 has no implications a priori for the sign-rank of $\mathbf{AC}^0$. Instead, our proofs of Theorems 5.2 and 5.3 are based on a stronger approximation-theoretic quantity that we call γ-*smooth threshold degree*. Formally, the γ-smooth threshold degree of a Boolean function $f\colon X \to \{-1, 1\}$ is the largest d for which there is a nonzero function $\phi\colon X \to \mathbb{R}$ with the following two properties: $\phi(x)f(x) \geq \gamma \cdot \|\phi\|_1/|X|$ for all $x \in X$; and $\langle \phi, p \rangle = 0$ for every polynomial of degree less than d. Taking $\gamma = 0$ in this formalism, one recovers the standard dual characterization of the threshold degree of f. In particular, threshold degree is synonymous with 0-smooth threshold degree. The general case of γ-smooth threshold degree for $\gamma > 0$ requires threshold degree witnesses ϕ that are *min-smooth*, in that the absolute value of ϕ at any given point is at least a γ fraction of the average absolute value of ϕ over all points.

The substantial advantage of *smooth* threshold degree is that it has immediate sign-rank implications. Specifically, any lower bound of d on the $2^{-O(d)}$-smooth threshold degree can be transformed efficiently and in a black-box manner into a sign-rank lower bound of $2^{\Omega(d)}$, using a combination of the pattern matrix method [**114, 116**] and Forster's spectral lower bound on sign-rank [**53, 54**]. Accordingly, we obtain Theorems 5.2 and 5.3 by proving an $\Omega(n^{1-\varepsilon})$ lower bound on the $2^{-n^{1-\varepsilon}}$-smooth threshold degree of $\mathbf{AC}^0$, for any constant $\varepsilon > 0$. At the core of this result is an amplification theorem for smooth threshold degree, whose repeated application makes it possible to prove arbitrarily strong lower bounds for $\mathbf{AC}^0$. Amplifying smooth threshold degree is a complex juggling act due to the presence of two parameters—degree and smoothness—that must evolve in coordinated fashion. The approach of Theorem 5.1 is not useful here because the threshold degree witnesses that arise from the proof of Theorem 5.1 are highly nonsmooth.

When amplifying the threshold degree of a function f as in the proof of Theorem 5.1, two phenomena adversely affect the smoothness parameter. The first is block-composition itself as a composition technique, which in the regime of interest to us transforms *every* threshold degree witness for f into a hopelessly nonsmooth witness for the composed function. The other culprit is the input compression step, which re-encodes the input and thereby affects the smoothness in ways that are hard to control. To overcome these difficulties, we develop a novel approach unrelated to our proof of Theorem 5.1.

Central to our work is an analytic property that we call *local smoothness.* Formally, let $\Phi\colon \mathbb{N}^n \to \mathbb{R}$ be a function of interest. For a subset $X \subseteq \mathbb{N}^n$ and a real number $K \geq 1$, we say that Φ is *K-smooth on X* if $|\Phi(x)| \leq K^{|x-x'|}|\Phi(x')|$ for all $x, x' \in X$. Put another way, for any two points of X at ℓ_1 distance d, the corresponding values of Φ differ in magnitude by a factor of at most K^d. In and of itself, a locally smooth function Φ need not be min-smooth because for a pair of points that are far from each other, the corresponding Φ-values can differ by many orders of magnitude. However, locally smooth functions exhibit extraordinary plasticity. Specifically, we show how to modify a locally smooth function's metric properties—such as its support or the distribution of its ℓ_1 mass—without the change being detectable by low-degree polynomials. This apparatus makes it possible to restore min-smoothness to the dual object Φ that results from the block-composition step and preserve that min-smoothness throughout the input compression step, eliminating the two obstacles to min-smoothness in the earlier proof of Theorem 5.1. The block-composition step here uses a *locally smooth* witness for the threshold degree of MP_m, which needs to be built from scratch and is quite different from the witness in the proof of Theorem 5.1.

Our described approach is quite different from previous work on the sign-rank of constant-depth circuits [**106, 37, 39**]. The analytic notion in those earlier papers is weaker than γ-smooth threshold degree and in particular allows the dual object to be *arbitrary* on a γ fraction of the inputs. This weaker property is acceptable when the main result is proved in one shot, with a closed-form construction of the dual object.

By contrast, we must construct dual objects iteratively, with each iteration increasing the degree parameter and proportionately decreasing the smoothness parameter. This iterative process requires that the dual object in each iteration be min-smooth on the entire domain. Perhaps unexpectedly, we find γ-smooth threshold degree easier to work with than the weaker notion in previous work [**106, 37, 39**]. In particular, we are able to give a new and short proof of the $\exp(\Omega(n^{1/3}))$ lower bound on the sign-rank of $\mathbf{AC}^0$, originally obtained by Razborov and Sherstov [**106**] with a much more complicated approach. The new proof can be found in Section 5.5.1, where it serves as a prelude to our main result on the sign-rank of $\mathbf{AC}^0$.

5.2. Preliminaries

5.2.1. Products. For a set X, we let $\mathbb{R}^X$ denote the linear space of real-valued functions on X. The *tensor product* of $f \in \mathbb{R}^X$ and $g \in \mathbb{R}^Y$ is denoted $f \otimes g \in \mathbb{R}^{X \times Y}$ and given by $(f \otimes g)(x, y) = f(x)g(y)$. The tensor product $f \otimes f \otimes \cdots \otimes f$ (n times) is abbreviated $f^{\otimes n}$. For a subset $S \subseteq \{1, 2, \ldots, n\}$ and a function $f\colon X \to \mathbb{R}$, we define $f^{\otimes S}\colon X^n \to \mathbb{R}$ by $f^{\otimes S}(x_1, x_2, \ldots, x_n) = \prod_{i \in S} f(x_i)$. As extremal cases, we have $f^{\otimes \varnothing} \equiv 1$ and $f^{\otimes \{1,2,\ldots,n\}} = f^{\otimes n}$. Tensor product notation generalizes naturally to *sets* of functions: $F \otimes G = \{f \otimes g : f \in F, g \in G\}$ and $F^{\otimes n} = \{f_1 \otimes f_2 \otimes \cdots \otimes f_n : f_1, f_2, \ldots, f_n \in F\}$. A *conical combination* of $f_1, f_2, \ldots, f_k \in \mathbb{R}^X$ is any function of the form $\lambda_1 f_1 + \lambda_2 f_2 + \cdots + \lambda_k f_k$, where $\lambda_1, \lambda_2, \ldots, \lambda_k$ are nonnegative reals. A *convex combination* of $f_1, f_2, \ldots, f_k \in \mathbb{R}^X$ is any function $\lambda_1 f_1 + \lambda_2 f_2 + \cdots + \lambda_k f_k$, where $\lambda_1, \lambda_2, \ldots, \lambda_k$ are nonnegative reals that sum to 1. The *conical hull* of $F \subseteq \mathbb{R}^X$, denoted $\operatorname{cone} F$, is the set of all conical combinations of functions in F. The *convex hull*, denoted $\operatorname{conv} F$, is defined analogously as the set of all convex combinations of functions in F. For any set of functions $F \subseteq \mathbb{R}^X$, we have

$$(\operatorname{conv} F)^{\otimes n} \subseteq \operatorname{conv}(F^{\otimes n}). \tag{5.2.1}$$

Throughout this manuscript, we view probability distributions as real functions. This convention makes available the shorthands introduced above. In particular, for probability distributions μ and λ, the symbol $\operatorname{supp} \mu$ denotes the support of μ, and $\mu \otimes \lambda$ denotes the probability distribution given by $(\mu \otimes \lambda)(x, y) = \mu(x)\lambda(y)$. If μ is a probability distribution on X, we consider μ to be defined also on any superset of X with the understanding that $\mu = 0$ outside X. We let $\mathfrak{D}(X)$ denote the family of all finitely supported probability distributions on X. Most of this chapter is concerned with the distribution family $\mathfrak{D}(\mathbb{N}^n)$ and its subfamilies, each of which we denote with a Fraktur letter. For any sets $X \subseteq \mathbb{N}^n$ and $W \subseteq \mathbb{R}$, we define

$$X|_W = \{x \in X : |x| \in W\}.$$

In the case of a one-element set $W = \{w\}$, we further shorten $X|_{\{w\}}$ to $X|_w$. To illustrate, $\mathbb{N}^n|_{\leq w}$ denotes the set of vectors whose components are natural numbers and sum to at most w, whereas $\{0,1\}^n|_w$ denotes the set of Boolean strings of length n and Hamming weight exactly w. For a function $f\colon X \to \mathbb{R}$ on a subset $X \subseteq \mathbb{N}^n$, we let $f|_W$ denote the restriction of f to $X|_W$. A typical use of this notation would be $f|_{\leq w}$ for some real number w.

5.2.2. Orthogonal content. For a multivariate real polynomial $p\colon \mathbb{R}^n \to \mathbb{R}$, we let $\deg p$ denote the total degree of p, i.e., the largest degree of any monomial of p. We use the terms *degree* and *total degree* interchangeably in this chapter. It will be convenient to define the degree of the zero polynomial by $\deg 0 = -\infty$. For a real-valued function ϕ supported on a finite subset of $\mathbb{R}^n$, we define the *orthogonal content of* ϕ, denoted $\operatorname{orth} \phi$, to be the minimum degree of a real polynomial p for which

$\langle \phi, p\rangle \neq 0$. We adopt the convention that $\operatorname{orth}\phi = \infty$ if no such polynomial exists. It is clear that $\operatorname{orth}\phi \in \mathbb{N}\cup\{\infty\}$, with the extremal cases $\operatorname{orth}\phi = 0 \;\Leftrightarrow\; \langle\phi, 1\rangle \neq 0$ and $\operatorname{orth}\phi = \infty \;\Leftrightarrow\; \phi = 0$. This gives us a simpler definition of the threshold degree and the smooth threshold degree for any $f : \{0,1\}^n \to \mathbb{R}$, in view of Equation (2.6.1)-(2.6.2),

$$\begin{aligned}
\deg_\pm(f) &= \max_{\mu\in\mathfrak{D}(X)} \operatorname{orth}((-1)^f\cdot\mu),\\
\deg_\pm(f,\gamma) &= \max_{\substack{\mu\in\mathfrak{D}(X):\\ \mu\geq\gamma/|X| \text{ on } X}} \operatorname{orth}((-1)^f\cdot\mu).
\end{aligned}$$

Our next three results record additional facts about orthogonal content.

PROPOSITION 5.7. *Let X and Y be nonempty finite subsets of Euclidean space. Then:*

(i) $\operatorname{orth}(\phi+\psi) \geq \min\{\operatorname{orth}\phi, \operatorname{orth}\psi\}$ for all $\phi,\psi\colon X\to\mathbb{R}$;

(ii) $\operatorname{orth}(\phi\otimes\psi) = \operatorname{orth}(\phi)+\operatorname{orth}(\psi)$ for all $\phi\colon X\to\mathbb{R}$ and $\psi\colon Y\to\mathbb{R}$;

(iii) $\operatorname{orth}(\phi^{\otimes n}-\psi^{\otimes n}) \geq \operatorname{orth}(\phi-\psi)$ for all $\phi,\psi\colon X\to\mathbb{R}$ and all $n\geq 1$.

Proof. Item (i) is immediate, as is the upper bound in (ii). For the lower bound in (ii), simply note that the linearity of inner product makes it possible to restrict attention to factored polynomials $p(x)q(y)$, where p and q are polynomials on X and Y, respectively. For (iii), use a telescoping sum to write

$$\begin{aligned}
\phi^{\otimes n}-\psi^{\otimes n} &= \sum_{i=0}^{n-1}(\phi^{\otimes(n-i)}\otimes\psi^{\otimes i} - \phi^{\otimes(n-i-1)}\otimes\psi^{\otimes(i+1)})\\
&= \sum_{i=0}^{n-1}\phi^{\otimes(n-i-1)}\otimes(\phi-\psi)\otimes\psi^{\otimes i}.
\end{aligned}$$

By (ii), each term in the final expression has orthogonal content at least $\operatorname{orth}(\phi-\psi)$. By (i), then, the sum has orthogonal content at least $\operatorname{orth}(\phi-\psi)$ as well. □

Proposition 5.8. *Let $\phi_0, \phi_1 \colon X \to \mathbb{R}$ be given functions on a finite subset X of Euclidean space. Then for every polynomial $p \colon X^n \to \mathbb{R}$, the mapping $z \mapsto \langle \bigotimes_{i=1}^n \phi_{z_i}, p\rangle$ is a polynomial on $\{0,1\}^n$ of degree at most $(\deg p)/\operatorname{orth}(\phi_1 - \phi_0)$.*

Proof. We may assume that $\operatorname{orth}(\phi_1 - \phi_0) > 0$ since the proposition holds trivially otherwise. By linearity, it suffices to consider factored polynomials $p(x_1, \dots, x_n) = \prod_{i=1}^n p_i(x_i)$, where each p_i is a nonzero polynomial on X. In this setting, we have

$$\left\langle \bigotimes_{i=1}^n \phi_{z_i}, p \right\rangle = \prod_{i=1}^n \langle \phi_{z_i}, p_i \rangle . \tag{5.2.2}$$

By definition, $\langle \phi_0, p_i\rangle = \langle \phi_1, p_i\rangle$ for any index i with $\deg p_i < \operatorname{orth}(\phi_1 - \phi_0)$. As a result, such indices do not contribute to the degree of the right-hand side of (5.2.2) as a function of z. The contribution of any other index to the degree is clearly at most 1. Summarizing, the right-hand side of (5.2.2) is a polynomial in $z \in \{0,1\}^n$ of degree at most $|\{i : \deg p_i \geq \operatorname{orth}(\phi_1 - \phi_0)\}| \leq (\deg p)/\operatorname{orth}(\phi_1 - \phi_0)$. □

Corollary 5.9. *Let X be a finite subset of Euclidean space. Then for any functions $\phi_0, \phi_1 \colon X \to \mathbb{R}$ and $\psi \colon \{0,1\}^n \to \mathbb{R}$,*

$$\operatorname{orth}\left(\sum_{z \in \{0,1\}^n} \psi(z) \bigotimes_{i=1}^n \phi_{z_i} \right) \geq \operatorname{orth}(\psi) \cdot \operatorname{orth}(\phi_1 - \phi_0).$$

Proof. We may assume that $\operatorname{orth}(\psi) \cdot \operatorname{orth}(\phi_1 - \phi_0) > 0$ since the claim holds trivially otherwise. Fix a polynomial any polynomial P of degree less than $\operatorname{orth}(\psi) \cdot \operatorname{orth}(\phi_1 - \phi_0)$. The linearity of inner product leads to

$$\left\langle \sum_{z \in \{0,1\}^n} \psi(z) \bigotimes_{i=1}^n \phi_{z_i}, P \right\rangle = \sum_{z \in \{0,1\}^n} \psi(z) \left\langle \bigotimes_{i=1}^n \phi_{z_i}, P \right\rangle .$$

By Proposition 5.8, the right-hand side is the inner product of ψ with a polynomial of degree less than $\operatorname{orth} \psi$ and is therefore zero. □

Observe that Corollary 5.9 gives an alternate proof of Proposition 5.7(iii). Our next proposition uses orthogonal content to give a useful criterion for a real-valued function to be a probability distribution.

PROPOSITION 5.10. *Let Λ be a probability distribution on a finite subset X of Euclidean space. Let $\tilde{\Lambda}\colon X \to \mathbb{R}$ be given with $\tilde{\Lambda} \geq 0$ and $\operatorname{orth}(\Lambda - \tilde{\Lambda}) > 0$. Then $\tilde{\Lambda}$ is a probability distribution on X.*

Proof. By hypothesis, $\tilde{\Lambda}$ is a nonnegative function. Moreover, $\|\tilde{\Lambda}\|_1 = \langle \tilde{\Lambda}, 1\rangle = \langle \Lambda, 1\rangle - \langle \Lambda - \tilde{\Lambda}, 1\rangle = \langle \Lambda, 1\rangle = 1$, where the third step uses $\operatorname{orth}(\Lambda - \tilde{\Lambda}) > 0$. □

Consider the real vector space of functions $\{0,1\}^n \to \mathbb{R}$. The linear subspace of real polynomials on $\{0,1\}^n$ of degree at most d is easily seen to be $\operatorname{span}\{\chi_S : |S| \leq d\}$. Its orthogonal complement, $\operatorname{span}\{\chi_S\colon |S| > d\}$, is then the linear subspace of functions that have zero inner product with every polynomial of degree at most d. As a result, the orthogonal content of a nonzero function $\phi\colon \{0,1\}^n \to \mathbb{R}$ is given by

$$\operatorname{orth}\phi = \min\{|S| : \hat{\phi}(S) \neq 0\}, \qquad \phi \not\equiv 0. \tag{5.2.3}$$

5.2.3. Symmetrization. Let S_n denote the symmetric group on n elements. For a permutation $\sigma \in S_n$ and an arbitrary sequence $x = (x_1, x_2, \ldots, x_n)$, we adopt the shorthand $\sigma x = (x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(n)})$. A function $f(x_1, x_2, \ldots, x_n)$ is called *symmetric* if it is invariant under permutation of the input variables: $f(x_1, x_2, \ldots, x_n) = f(x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(n)})$ for all x and σ. Symmetric functions on $\{0,1\}^n$ are intimately related to univariate polynomials, as was first observed by Minsky and Papert in their *symmetrization argument* [**90**].

Proposition 5.11 (Minsky and Papert). *Let $p\colon \mathbb{R}^n \to \mathbb{R}$ be a given polynomial. Then the mapping*

$$t \mapsto \mathop{\mathbf{E}}_{x \in \{0,1\}^n|_t} p(x)$$

is a univariate polynomial on $\{0, 1, 2, \dots, n\}$ of degree at most $\deg p$.

Minsky and Papert's result generalizes to block-symmetric functions:

Proposition 5.12. *Let $n_1, \dots, n_k$ be positive integers. Let $p\colon \mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_k} \to \mathbb{R}$ be a given polynomial. Then the mapping*

$$(t_1, t_2, \dots, t_k) \mapsto \mathop{\mathbf{E}}_{x_1 \in \{0,1\}^{n_1}|_{t_1}} \mathop{\mathbf{E}}_{x_2 \in \{0,1\}^{n_2}|_{t_2}} \cdots \mathop{\mathbf{E}}_{x_k \in \{0,1\}^{n_k}|_{t_k}} p(x_1, x_2, \dots, x_k)$$

is a polynomial on $\{0, 1, \dots, n_1\} \times \{0, 1, \dots, n_2\} \times \cdots \times \{0, 1, \dots, n_k\}$ of degree at most $\deg p$.

Proposition 5.12 follows in a straightforward manner from Proposition 5.11 by induction on the number of blocks k, as pointed out in [**106**, Proposition 2.3]. The next result is yet another generalization of Minsky and Papert's symmetrization technique, this time to the setting when $x_1, x_2, \dots, x_n$ are vectors rather than bits.

Proposition 5.13. *Let $p\colon (\mathbb{R}^m)^n \to \mathbb{R}$ be a polynomial of degree d. Then there is a polynomial $p^*\colon \mathbb{R}^n \to \mathbb{R}$ of degree at most d such that for all $x_1, x_2, \dots, x_n \in \{e_1, e_2, \dots, e_m, 0^m\}$,*

$$\mathop{\mathbf{E}}_{\sigma \in S_n} p(x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(n)}) = p^*(x_1 + x_2 + \cdots + x_n).$$

Proof. We closely follow an argument due to Ambainis [**8**, Lemma 3.4], who proved a related result. Since the components of $x_1, x_2, \dots, x_n$ are Boolean-valued, we have $x_{i,j} = x_{i,j}^2 = x_{i,j}^3 = \cdots$ and therefore we may assume that p is multilinear. By linearity, it further suffices to consider the case when p is a single monomial:

$$p(x_1, x_2, \ldots, x_n) = \prod_{j=1}^{m} \prod_{i \in S_j} x_{i,j} \tag{5.2.4}$$

for some sets $S_1, S_2, \ldots, S_m \subseteq \{1, 2, \ldots, n\}$ with $\sum_{j=1}^{m} |S_j| \leq d$. If some pair of sets $S_j, S_{j'}$ with $j \neq j'$ have nonempty intersection, then the right-hand side of (5.2.4) contains a product of the form $x_{i,j} x_{i,j'}$ for some i and thus $p \equiv 0$ on the domain in question. As a result, the proposition holds with $p^* = 0$. In the complementary case when $S_1, S_2, \ldots, S_m$ are pairwise disjoint, we calculate

$$\begin{aligned}
&\mathop{\mathbf{E}}_{\sigma \in S_n} p(x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(n)}) \\
&\qquad = \prod_{j=1}^{m} \mathop{\mathbf{E}}_{\sigma \in S_n} \left[\prod_{i \in S_j} x_{\sigma(i),j} \,\middle|\, \prod_{i \in S_{j'}} x_{\sigma(i),j'} = 1 \text{ for all } j' < j \right] \\
&\qquad = \prod_{j=1}^{m} \binom{x_{1,j} + x_{2,j} + \cdots + x_{n,j}}{|S_j|} \binom{n - |S_1| - |S_2| - \cdots - |S_{j-1}|}{|S_j|}^{-1}.
\end{aligned}$$

Expanding out the binomial coefficients shows that the final expression is an m-variate polynomial whose argument is the vector sum $x_1 + x_2 + \cdots + x_n \in \mathbb{R}^m$. Moreover, the degree of this polynomial is $\sum |S_j| \leq d$. □

Corollary 5.14. *Let $p\colon (\mathbb{R}^m)^n \to \mathbb{R}$ be a polynomial of degree d. Then the mapping*

$$v \mapsto \mathop{\mathbf{E}}_{\substack{x \in \{0^m, e_1, e_2, \ldots, e_m\}^n: \\ x_1 + x_2 + \cdots + x_n = v}} p \tag{5.2.5}$$

is a polynomial on $\mathbb{N}^m|_{\leq n}$ of degree at most $\deg p$.

Minsky and Papert's symmetrization corresponds to $m = 1$ in Corollary 5.14.

Proof of Corollary 5.14. Let $v \in \mathbb{N}^m|_{\leq n}$ be given. Then all representations $v = x_1 + x_2 + \cdots + x_n$ with $x_1, x_2, \ldots, x_n \in \{0^m, e_1, e_2, \ldots, e_m\}$ are the same up to the order of the summands. As a result, (5.2.5) is the same mapping as

$$v \mapsto \mathop{\mathbf{E}}_{\sigma \in S_n} p(\sigma(\underbrace{e_1, \ldots, e_1}_{v_1}, \underbrace{e_2, \ldots, e_2}_{v_2}, \ldots, \underbrace{e_m, \ldots, e_m}_{v_m}, \underbrace{0^m, 0^m \ldots, 0^m}_{n-v_1-\cdots-v_m})),$$

which by Proposition 5.13 is a polynomial in

$$\underbrace{e_1 + \cdots + e_1}_{v_1} + \underbrace{e_2 + \cdots + e_2}_{v_2} + \cdots + \underbrace{e_m + \cdots + e_m}_{v_m} + \underbrace{0^m + \cdots + 0^m}_{n-v_1-\cdots-v_m} = v$$

of degree at most $\deg p$. □

Analogous to symmetrized polynomials, it will be also helplful to work with symmetrized versions of Boolean functions. We define $\mathrm{AND}_n^*, \mathrm{OR}_n^* \colon \{0, 1, 2, \ldots, n\} \to \{0, 1\}$ by

$$\mathrm{AND}_n^*(t) = \begin{cases} 1 & \text{if } t = n, \\ 0 & \text{otherwise,} \end{cases} \qquad \mathrm{OR}_n^*(t) = \begin{cases} 0 & \text{if } t = 0, \\ 1 & \text{otherwise.} \end{cases}$$

The symmetrized variant of the Minsky–Papert function is $\mathrm{MP}_{m,r}^* = \mathrm{AND}_m \circ \mathrm{OR}_r^*$.

5.3. Auxiliary results

In this section, we collect a number of supporting results on approximate degree that have appeared in one form or another in previous work. For the reader's convenience, we provide self-contained proofs whenever the precise formulation that we need departs from published work.

5.3.1. Basic dual objects.

As described in the introduction, we prove our main results constructively, by building explicit dual objects that witness the corresponding lower bounds. An important tool in this process is the following lemma due to Razborov and Sherstov [**106**]. Informally, it is used to adjust a dual object's metric properties while preserving its orthogonality to low-degree polynomials. The lemma plays a basic role in several recent papers [**106, 38, 33**] as well as our work.

LEMMA 5.15 (Razborov and Sherstov). *Fix integers d and n, where $0 \le d < n$. Then there is an (explicitly given) function $\zeta\colon \{0,1\}^n \to \mathbb{R}$ such that*

$$\operatorname{supp}\zeta \subseteq \{0,1\}^n|_{\le d} \cup \{1^n\},$$

$$\zeta(1^n) = 1,$$

$$\|\zeta\|_1 \le 1 + 2^d \binom{n}{d},$$

$$\operatorname{orth}\zeta > d.$$

In more detail, this result corresponds to taking $k = d$ and $\zeta = (-1)^n g$ in the proof of Lemma 3.2 of [**106**]. We will need the following symmetrized version of Lemma 5.15.

LEMMA 5.16. *Fix a point $u \in \mathbb{N}^n$ and a natural number $d < |u|$. Then there is $\zeta_u\colon \mathbb{N}^n \to \mathbb{R}$ such that*

$$\operatorname{supp}\zeta_u \subseteq \{u\} \cup \{v \in \mathbb{N}^n : v \le u \text{ and } |v| \le d\}, \tag{5.3.1}$$

$$\zeta_u(u) = 1, \tag{5.3.2}$$

$$\|\zeta_u\|_1 \le 1 + 2^d \binom{|u|}{d}, \tag{5.3.3}$$

$$\operatorname{orth}\zeta_u > d. \tag{5.3.4}$$

Proof. Lemma 5.15 gives a function $\zeta\colon \{0,1\}^{|u|} \to \mathbb{R}$ such that

$$\operatorname{supp}\zeta \subseteq \{0,1\}^{|u|}|_{\leq d} \cup \{1^{|u|}\}, \tag{5.3.5}$$

$$\zeta(1^{|u|}) = 1, \tag{5.3.6}$$

$$\|\zeta\|_1 \leq 1 + 2^d \binom{|u|}{d}, \tag{5.3.7}$$

$$\operatorname{orth}\zeta > d. \tag{5.3.8}$$

Now define $\zeta_u\colon \mathbb{N}^n \to \mathbb{R}$ by

$$\zeta_u(v) = \sum_{x_1 \in \{0,1\}^{|u_1|}|_{|v_1|}} \cdots \sum_{x_n \in \{0,1\}^{|u_n|}|_{|v_n|}} \zeta(x_1 \ldots x_n).$$

Then (5.3.1)–(5.3.3) are immediate from (5.3.5)–(5.3.7), respectively. To verify the remaining property (5.3.4), fix a polynomial $p\colon \mathbb{R}^n \to \mathbb{R}$ of degree at most d. Then

$$\begin{aligned}
\langle \zeta_u, p\rangle &= \sum_{v : v \leq u} \left(\sum_{x_1 \in \{0,1\}^{|u_1|}|_{|v_1|}} \cdots \sum_{x_n \in \{0,1\}^{|u_n|}|_{|v_n|}} \zeta(x_1 \ldots x_n) \right) p(v_1, \ldots, v_n) \\
&= \sum_{v : v \leq u} \left(\sum_{x_1 \in \{0,1\}^{|u_1|}|_{|v_1|}} \cdots \sum_{x_n \in \{0,1\}^{|u_n|}|_{|v_n|}} \zeta(x_1 \ldots x_n) p(|x_1|, \ldots, |x_n|) \right) \\
&= \sum_{x_1 \in \{0,1\}^{|u_1|}} \cdots \sum_{x_n \in \{0,1\}^{|u_n|}} \zeta(x_1 \ldots x_n) p(|x_1|, \ldots, |x_n|) \\
&= 0,
\end{aligned}$$

where the last step uses (5.3.8). □

When constructing a dual polynomial for a complicated constant-depth circuit, it is natural to start with a dual polynomial for the OR function or, equivalently, its counterpart AND. The first such dual polynomial was constructed by Špalek [**137**], with many refinements and generalizations [**34, 122, 124, 38, 33**] obtained in follow-up work. We augment this line of work with yet another construction, which delivers the exact combination of analytic and metric properties that we need.

THEOREM 5.17. *Let $0 < \varepsilon < 1$ be given. Then for some constants $c', c'' \in (0,1)$ and all integers $N \geq n \geq 1$, there is an (explicitly given) function $\psi\colon \{0,1,2,\ldots,N\} \to \mathbb{R}$ such that*

$$\begin{aligned}
&\psi(0) > \frac{1-\varepsilon}{2}, \\
&\|\psi\|_1 = 1, \\
&\operatorname{orth}\psi \geq c'\sqrt{n}, \\
&\operatorname{sgn}\psi(t) = (-1)^t, && t = 0,1,2,\ldots,N, \\
&|\psi(t)| \in \left[\frac{c'}{(t+1)^2\, 2^{c''t/\sqrt{n}}},\ \frac{1}{c'(t+1)^2\, 2^{c''t/\sqrt{n}}}\right], && t = 0,1,2,\ldots,N.
\end{aligned}$$

A self-contained proof of Theorem 5.17 is available in Appendix 5.6.

5.3.2. Dominant components. We now recall a lemma due to Bun and Thaler [**38**] that serves to identify the dominant components of a vector. Its primary use [**38, 33**] is to prove concentration-of-measure results for product distributions on $\mathbb{N}^n$.

LEMMA 5.18 (Bun and Thaler). *Let $v \in \mathbb{R}^n$ be given, $v \neq 0^n$. Then there is $S \subseteq \{1,2,\ldots,n\}$ such that*

$$\begin{aligned}
&|S| \geq \frac{\|v\|_1}{2\|v\|_\infty}, \\
&|S| \min_{i\in S} |v_i| \geq \frac{\|v\|_1}{2(1+\ln n)}.
\end{aligned}$$

Proof (adapted from [**38**]). By renumbering the indices if necessary, we may assume that $|v_1| \geq |v_2| \geq \cdots \geq |v_n| \geq 0$. For the sake of contradiction, suppose that no such set S exists. Then

$$|v_i| < \frac{1}{i} \cdot \frac{\|v\|_1}{2(1 + \ln n)}$$

for every index $i \geq \frac{\|v\|_1}{2\|v\|_\infty}$. As a result,

$$\begin{aligned}
\|v\|_1 &= \sum_{i < \frac{\|v\|_1}{2\|v\|_\infty}} |v_i| + \sum_{i = \left\lceil \frac{\|v\|_1}{2\|v\|_\infty} \right\rceil}^{n} |v_i| \\
&\leq \sum_{i < \frac{\|v\|_1}{2\|v\|_\infty}} \|v\|_\infty + \sum_{i = \left\lceil \frac{\|v\|_1}{2\|v\|_\infty} \right\rceil}^{n} \frac{1}{i} \cdot \frac{\|v\|_1}{2(1 + \ln n)} \\
&< \frac{\|v\|_1}{2} + \frac{\|v\|_1}{2(1 + \ln n)} \sum_{i=1}^{n} \frac{1}{i} \\
&\leq \|v\|_1,
\end{aligned}$$

where the final step uses

$$\sum_{i=1}^{n} \frac{1}{i} = 1 + \sum_{i=2}^{n} \frac{1}{i} \leq 1 + \int_1^n \frac{di}{i} = 1 + \ln n.$$

We have arrived at $\|v\|_1 < \|v\|_1$, a contradiction. □

We will need a slightly more general statement, which can be thought of as an extremal analogue of Lemma 5.18.

Lemma 5.19. *Fix $\theta > 0$ and let $v \in \mathbb{R}^n$ be an arbitrary vector with $\|v\|_1 \geq \theta$. Then there is $S \subseteq \{1, 2, \ldots, n\}$ such that*

$$|S| \geq \frac{\|v\|_1}{2\|v\|_\infty}, \tag{5.3.9}$$

$$\min_{i \in S} |v_i| \geq \frac{1}{|S|} \cdot \frac{\theta}{2(1 + \ln n)}, \tag{5.3.10}$$

$$\sum_{i \notin S} |v_i| < \theta. \tag{5.3.11}$$

Proof. Fix n, v, and θ for the remainder of the proof. We will refer to a subset $S \subseteq \{1, 2, \ldots, n\}$ as *regular* if S satisfies (5.3.9) and (5.3.10). Lemma 5.18 along with $\|v\|_1 \geq \theta$ ensures the existence of at least one regular set. Now, let S be a *maximal* regular set. For the sake of contradiction, suppose that (5.3.11) fails. Applying Lemma 5.18 to $v|_{\overline{S}}$ produces a nonempty set $T \subseteq \overline{S}$ with

$$\min_{i \in T} |v_i| \geq \frac{1}{|T|} \cdot \frac{\theta}{2(1 + \ln n)}.$$

But then $S \cup T$ is regular, contradicting the maximality of S. □

Lemmas 5.18 and 5.19 imply the following concentration-of-measure result for product distributions on $\mathbb{N}^n$, due to Bun and Thaler [**38**].

Lemma 5.20 (Bun and Thaler). *Let $\lambda_1, \lambda_2, \ldots, \lambda_n \in \mathfrak{D}(\mathbb{N})$ be given with*

$$\lambda_i(t) \leq \frac{C\alpha^t}{(t+1)^2}, \qquad t \in \mathbb{N}, \tag{5.3.12}$$

where $C \geq 0$ and $0 \leq \alpha \leq 1$. Then for all $\theta \geq 8Cen(1 + \ln n)$,

$$\mathop{\mathbf{P}}_{v \sim \lambda_1 \times \lambda_2 \times \cdots \times \lambda_n} [\|v\|_1 \geq \theta] \leq \alpha^{\theta/2}.$$

Proof (adapted from [**38**]). For a nonempty subset $S \subseteq \{1, 2, \ldots, n\}$ and a vector $v \in \mathbb{N}^n$, we say that *v is S-heavy* if the following conditions are simultaneously satisfied:

$$|v_i| \geq \frac{1}{|S|} \cdot \frac{\theta}{4(1+\ln n)}, \qquad\qquad i \in S, \tag{5.3.13}$$

$$\sum_{i\in S} |v_i| > \frac{\theta}{2}. \tag{5.3.14}$$

Now, consider a random vector $v \in \mathbb{N}^n$ distributed according to $\lambda_1 \times \lambda_2 \times \cdots \times \lambda_n$. We have

$$\begin{aligned}
\mathop{\mathbf{P}}_{v}[\|v\|_1 \geq \theta] &\leq \mathop{\mathbf{P}}_{v}[v \text{ is } S\text{-heavy for some nonempty } S \neq \varnothing] \\
&\leq \sum_{\substack{S\subseteq\{1,2,\ldots,n\}\\ S\neq\varnothing}} \mathop{\mathbf{P}}_{v}[v \text{ is } S\text{-heavy}] \\
&\leq \sum_{\substack{S\subseteq\{1,2,\ldots,n\}\\ S\neq\varnothing}} \alpha^{\theta/2} \left(\sum_{t \geq \frac{1}{|S|}\cdot\frac{\theta}{4(1+\ln n)}} \frac{C}{(t+1)^2} \right)^{|S|} \\
&\leq \sum_{\substack{S\subseteq\{1,2,\ldots,n\}\\ S\neq\varnothing}} \alpha^{\theta/2} \left(C \int_{\frac{1}{|S|}\cdot\frac{\theta}{4(1+\ln n)}}^{\infty} \frac{dt}{t^2} \right)^{|S|} \\
&= \sum_{\substack{S\subseteq\{1,2,\ldots,n\}\\ S\neq\varnothing}} \alpha^{\theta/2} \left(\frac{C|S| \cdot 4(1+\ln n)}{\theta} \right)^{|S|} \\
&= \sum_{s=1}^{n} \binom{n}{s} \cdot \alpha^{\theta/2} \left(\frac{Cs \cdot 4(1+\ln n)}{\theta} \right)^{s} \\
&\leq \sum_{s=1}^{n} \alpha^{\theta/2} \left(\frac{en}{s} \cdot \frac{Cs \cdot 4(1+\ln n)}{\theta} \right)^{s} \\
&\leq \alpha^{\theta/2},
\end{aligned}$$

where the first inequality holds by Lemma 5.19; the second step applies the union bound; the third step uses $0 \leq \alpha \leq 1$ and the upper bound (5.3.12) for the λ_i; and the last two steps use (2.1.1) and the hypothesis that $\theta \geq 8Cen(1+\ln n)$, respectively. □

5.3.3. Input transformation. We work almost exclusively with Boolean functions on $\mathbb{N}^n|_{\leq\theta}$, where the dimension parameter n is polynomially larger than the Hamming weight parameter θ. This choice of domain is admittedly unusual but greatly simplifies the analysis. Fortunately, approximation-theoretic results obtained in this setting carry over in a blackbox manner to the hypercube. In more detail, we will now prove that every function on $\mathbb{N}^n|_{\leq\theta}$ can be transformed into a function on $O(\theta\log n)$ Boolean variables with similar approximation-theoretic properties. Analogous input transformations, with similar proofs, have been used in previous work to translate results from $\{0,1\}^n|_\theta$ or $\{0,1\}^n|_{\leq\theta}$ to the hypercube setting [**38, 33**]. The presentation below seems more economical than previous treatments.

Recall that $e_1, e_2, \ldots, e_n$ denote the standard basis for $\mathbb{R}^n$. The following encoding lemma was proved in [**124**, Lemma 3.1].

Lemma 5.21 (Sherstov). *Let $n \geq 1$ be a given integer. Then there is a surjection $g\colon \{0,1\}^{6\lceil\log(n+1)\rceil} \to \{0^n, e_1, e_2, \ldots, e_n\}$ such that*

$$\mathop{\mathbf{E}}_{g^{-1}(0^n)} p = \mathop{\mathbf{E}}_{g^{-1}(e_1)} p = \mathop{\mathbf{E}}_{g^{-1}(e_2)} p = \cdots = \mathop{\mathbf{E}}_{g^{-1}(e_n)} p$$

for every polynomial p of degree at most $\lceil\log(n+1)\rceil$. Moreover, g can be constructed deterministically in time polynomial in n.

Observe that the points $0^n, e_1, e_2, \ldots, e_n$ in this lemma act simply as labels and can be replaced with any other tuple of $n+1$ distinct points. Indeed, this result was originally stated in [**124**] for a different choice of points. A tensor version of Lemma 5.21 is as follows.

Lemma 5.22. *Let $g\colon \{0,1\}^{6\lceil\log(n+1)\rceil} \to \{0^n, e_1, e_2, \ldots, e_n\}$ be as constructed in Lemma 5.21. Then for any integer $\theta \geq 1$ and for any polynomial $p\colon (\mathbb{R}^{6\lceil\log(n+1)\rceil})^\theta \to$*

$\mathbb{R}$, *the mapping*

$$(y_1, y_2, \ldots, y_\theta) \mapsto \mathop{\mathbf{E}}_{g^{-1}(y_1)\times g^{-1}(y_2)\times\cdots\times g^{-1}(y_\theta)} p$$

is a polynomial in $y \in \{0^n, e_1, e_2, \ldots, e_n\}^\theta$ *of degree at most* $(\deg p)/\lceil\log(n+1)+1\rceil$.

Proof. By linearity, it suffices to prove consider factored polynomials of the form $p(x_1, x_2, \ldots, x_\theta) = p_1(x_1)p_2(x_2)\cdots p_\theta(x_\theta)$, where $p_1, p_2, \ldots, p_\theta$ are real polynomials on $\{0,1\}^{6\lceil\log(n+1)\rceil}$. For such a polynomial, the defining equation simplifies to

$$\mathop{\mathbf{E}}_{g^{-1}(y_1)\times g^{-1}(y_2)\times\cdots\times g^{-1}(y_\theta)} p \;= \prod_{i=1}^{n} \mathop{\mathbf{E}}_{g^{-1}(y_i)} p_i. \tag{5.3.15}$$

We now examine the individual contributions of $p_1, p_2, \ldots, p_\theta$ to the degree of the right-hand side as a real polynomial in y. For any polynomial p_i of degree at most $\lceil\log(n+1)\rceil$, Lemma 5.21 ensures that the corresponding expectation $\mathbf{E}_{g^{-1}(y_i)} p_i$ is a constant independent of the input y_i. Thus, polynomials p_i of degree at most $\lceil\log(n+1)\rceil$ do not contribute to the degree of the right-hand side of (5.3.15). For the other polynomials p_i, the expectation $\mathbf{E}_{g^{-1}(y_i)} p_i$ is a linear polynomial in y_i, namely,

$$\begin{aligned}\mathop{\mathbf{E}}_{g^{-1}(y_i)} p_i = y_{i,1} \mathop{\mathbf{E}}_{g^{-1}(e_1)} p_i + y_{i,2} \mathop{\mathbf{E}}_{g^{-1}(e_2)} p_i + \cdots + y_{i,n} \mathop{\mathbf{E}}_{g^{-1}(e_n)} p_i \\ + \left(1 - \sum_{j=1}^{n} y_{i,j}\right) \mathop{\mathbf{E}}_{g^{-1}(0^n)} p_i,\end{aligned}$$

where we are crucially exploiting the fact that $y_i \in \{0^n, e_1, e_2, \ldots, e_n\}$. Thus, polynomials p_i of degree greater than $\lceil\log(n+1)\rceil$ contribute at most 1 each to the degree. Summarizing, the right-hand side of (5.3.15) is a real polynomial in $y_1, y_2, \ldots, y_\theta$ of degree at most

$$|\{i : \deg p_i \geq \lceil\log(n+1)\rceil + 1\}| \leq \frac{\deg p}{\lceil\log(n+1)\rceil + 1}. \qquad \square$$

We have reached the claimed result on input transformation.

THEOREM 5.23. *Let $n, \theta \geq 1$ be given integers. Set $N = 6\lceil \log(n+1) \rceil \theta$. There is a surjection $G\colon \{0,1\}^N \to \mathbb{N}^n|_{\leq \theta}$ such that:*

(i) *for every polynomial $p\colon \mathbb{R}^N \to \mathbb{R}$, the mapping $v \mapsto \mathbf{E}_{G^{-1}(v)}\, p$ is a polynomial on $\mathbb{N}^n|_{\leq \theta}$ of degree at most $(\deg p)/\lceil \log(n+1) + 1 \rceil$;*

(ii) *for every coordinate $i = 1, 2, \ldots, n$, the mapping $x \mapsto \mathrm{OR}^*_\theta(G(x)_i)$ is computable by an explicitly given DNF formula with $O(\theta n^6)$ terms, each with at most $6\lceil \log(n+1) \rceil$ variables.*

Applying Theorem 5.23 to a function $f\colon \mathbb{N}^n|_{\leq \theta} \to \{0,1\}$ produces a composed function $f \circ G\colon \{0,1\}^{6\lceil \log(n+1)\rceil \theta} \to \{0,1\}$ in the hypercube setting. The theorem ensures that lower bounds for the pointwise approximation, or sign-representation, of f apply to $f \circ G$ as well. Moreover, the circuit complexity of $f \circ G$ is only slightly higher than that of f. This way, Theorem 5.23 efficiently transfers approximation-theoretic results from $\mathbb{N}^n|_{\leq \theta}$ (or any subset thereof, such as $\{0,1\}^n|_{\leq \theta}$ or $\mathbb{N}^n|_{\theta}$) to the traditional setting of the hypercube.

Proof of Theorem 5.23. Define $G\colon (\{0,1\}^{6\lceil \log(n+1)\rceil})^\theta \to \mathbb{N}^n|_{\leq \theta}$ by

$$G(x_1, x_2, \ldots, x_\theta) = g(x_1) + g(x_2) + \cdots + g(x_\theta),$$

where $g\colon \{0,1\}^{6\lceil \log(n+1)\rceil} \to \{0^n, e_1, e_2, \ldots, e_n\}$ is as constructed in Lemma 5.21. The surjectivity of G follows trivially from that of g. We proceed to verify the additional properties required of G.

(i) For $v \in \mathbb{N}^n|_{\leq \theta}$, we have the partition

$$G^{-1}(v) = \bigcup_{\substack{y \in \{0^n, e_1, e_2, \ldots, e_n\}^\theta: \\ y_1 + y_2 + \cdots + y_\theta = v}} g^{-1}(y_1) \times g^{-1}(y_2) \times \cdots \times g^{-1}(y_\theta). \tag{5.3.16}$$

All representations $v = y_1+y_2+\cdots+y_\theta$ with $y_1, y_2, \ldots, y_\theta \in \{0^n, e_1, e_2, \ldots, e_n\}$ are the same up to the order of the summands. As a result, each part $g^{-1}(y_1)\times g^{-1}(y_2)\times\cdots\times g^{-1}(y_\theta)$ in the partition on the right-hand side of (5.3.16) has the same cardinality. We conclude that for any given polynomial p,

$$\mathop{\mathbf{E}}_{G^{-1}(v)} p = \mathop{\mathbf{E}}_{\substack{y\in\{0^n,e_1,e_2,\ldots,e_n\}^\theta:\\ y_1+y_2+\cdots+y_\theta=v}} \quad \mathop{\mathbf{E}}_{g^{-1}(y_1)\times g^{-1}(y_2)\times\cdots\times g^{-1}(y_\theta)} p. \tag{5.3.17}$$

Recall from Lemma 5.22 that the rightmost expectation in this equation is a polynomial in $y_1, y_2, \ldots, y_\theta \in \{0^n, e_1, e_2, \ldots, e_n\}$ of degree at most $(\deg p)/\lceil\log(n+1)+1\rceil$. As a result, Corollary 5.14 implies that the right-hand side of (5.3.17) is a polynomial in v of degree at most $(\deg p)/\lceil\log(n+1)+1\rceil$.

(ii) Fix an index i. Then

$$\mathrm{OR}^*_\theta(G(x)_i) = \bigvee_{j=1}^{\theta} \mathbf{I}[g(x_j) = e_i].$$

Each of the disjuncts on the right-hand side is a function of $6\lceil\log(n+1)\rceil$ Boolean variables. Therefore, $\mathrm{OR}^*_\theta(G(x)_i)$ is representable by a DNF formula with $O(\theta n^6)$ terms, each with at most $6\lceil\log(n+1)\rceil$ variables. □

5.4. The threshold degree of $\mathbf{AC}^0$

This section is devoted to our results on threshold degree. While we are mainly interested in the threshold degree of $\mathbf{AC}^0$, the techniques developed here apply to a much broader class of functions. Specifically, we prove an *amplification theorem* that takes an arbitrary function f and builds from it a function F with higher threshold degree. We give analogous amplification theorems for various other approximation-theoretic quantities. The transformation $f \mapsto F$ is efficient with regard to circuit depth and size and in particular preserves membership in $\mathbf{AC}^0$. To deduce our main results for

$\mathbf{AC}^0$, we start with a single-gate circuit and iteratively apply the amplification theorem to produce constant-depth circuits of higher and higher threshold degree. We develop this general machinery in Sections 5.4.1–5.4.3, followed by the application to $\mathbf{AC}^0$ in Section 5.4.5.

5.4.1. Shifting probability mass in product distributions. Consider a product distribution Λ on $\mathbb{N}^n$ whereby every component is concentrated near 0. The centerpiece of our work, presented here, is the construction of an associated probability distribution $\tilde{\Lambda}$ that is supported entirely on inputs of low weight and cannot be distinguished from Λ by a low-degree polynomial. More formally, define $\mathfrak{B}(r, c, \alpha)$ to be the family of probability distributions λ on $\mathbb{N}$ such that

$$\operatorname{supp} \lambda = \{0, 1, 2, \ldots, r'\}$$

for some nonnegative integer $r' \leq r$, and in addition

$$\frac{c^{t+1}}{(t+1)^2\, 2^{\alpha t}} \leq \lambda(t) \leq \frac{1}{c(t+1)^2\, 2^{\alpha t}}, \qquad t \in \operatorname{supp} \lambda. \tag{5.4.1}$$

Distributions in this family are subject to pointwise constraints, hence the symbol $\mathfrak{B}$ for "bounded." Our choice of bounding functions is motivated mainly by the metric properties of the dual polynomial for OR_n, constructed in Theorem 5.17.

In this notation, our analysis handles any distribution $\Lambda \in \mathfrak{B}(r, c, \alpha)^{\otimes n}$. It would be possible to generalize our work further, but the lower and upper bounds in (5.4.1) are already exponentially far apart and capture a much larger class of probability distributions than what we need for the applications to $\mathbf{AC}^0$. The precise statement of our result is as follows.

THEOREM 5.24. *Let $\Lambda \in \mathfrak{B}(r,c,\alpha)^{\otimes n}$ be given, for some integer $r \geq 0$ and reals $c > 0$ and $\alpha \geq 0$. Let d and θ be positive integers with*

$$\theta \geq 2d, \tag{5.4.2}$$

$$\theta \geq \frac{4en(1+\ln n)}{c^2}. \tag{5.4.3}$$

Then there is a function $\tilde{\Lambda}\colon \mathbb{N}^n \to \mathbb{R}$ such that

$$\operatorname{supp}\tilde{\Lambda} \subseteq (\operatorname{supp}\Lambda)|_{<2\theta}, \tag{5.4.4}$$

$$\operatorname{orth}(\Lambda - \tilde{\Lambda}) > d, \tag{5.4.5}$$

$$|\Lambda - \tilde{\Lambda}| \leq \left(\frac{8nr}{c}\right)^d 2^{-\lceil \theta/r \rceil - \alpha\lceil \theta/2 \rceil + 2}\,\Lambda \qquad \text{on } \operatorname{supp}\tilde{\Lambda}. \tag{5.4.6}$$

In general, the function $\tilde{\Lambda}$ constructed in Theorem 5.24 may not be a probability distribution. However, when θ is large enough relative to the other parameters, the pointwise property (5.4.6) forces $|\Lambda - \tilde{\Lambda}| \leq \Lambda$ and in particular $\tilde{\Lambda} \geq 0$. Since $\operatorname{orth}(\Lambda - \tilde{\Lambda}) > 0$ by construction, Proposition 5.10 guarantees that $\tilde{\Lambda}$ *is* a probability distribution in that case.

Proof of Theorem 5.24. For $c > 1$, we have $\mathfrak{B}(r,c,\alpha) = \varnothing$ and the theorem holds vacuously. Another degenerate possibility is $r = 0$, in which case Λ is the single-point distribution on 0^n, and therefore it suffices to take $\tilde{\Lambda} = \Lambda$. In what follows, we treat the general case when

$$c \in (0,1],$$

$$r \geq 1.$$

For every vector $v \in \mathbb{N}^n$ with $\|v\|_1 \geq \theta$, let $S(v) \subseteq \{1, 2, \ldots, n\}$ denote the corresponding subset identified by Lemma 5.19. To restate the lemma's guarantees,

$$|S(v)| \geq \frac{\theta}{r}, \qquad v \in (\operatorname{supp} \Lambda)|_{\geq 2\theta}, \tag{5.4.7}$$

$$\min_{i \in S(v)} v_i \geq \frac{\theta}{2|S(v)|(1 + \ln n)}, \qquad v \in (\operatorname{supp} \Lambda)|_{\geq 2\theta}, \tag{5.4.8}$$

$$\|v|_{\overline{S(v)}}\|_1 < \theta. \qquad v \in (\operatorname{supp} \Lambda)|_{\geq 2\theta}. \tag{5.4.9}$$

Property (5.4.9) implies that

$$\|v|_{S(v)}\|_1 > \theta, \qquad v \in (\operatorname{supp} \Lambda)|_{\geq 2\theta}, \tag{5.4.10}$$

and in particular

$$\|v|_{S(v)}\|_1 > d, \qquad v \in (\operatorname{supp} \Lambda)|_{\geq 2\theta}. \tag{5.4.11}$$

For each $i = 1, 2, \ldots, n$ and each $u \in \mathbb{N}^i|_{>d}$, Lemma 5.16 gives a function $\zeta_u \colon \mathbb{N}^i \to \mathbb{R}$ such that

$$\operatorname{supp} \zeta_u \subseteq \{u\} \cup \{v \in \mathbb{N}^i : v \leq u \text{ and } |v| \leq d\}, \tag{5.4.12}$$

$$\zeta_u(u) = 1, \tag{5.4.13}$$

$$\|\zeta_u\|_1 \leq 1 + 2^d \binom{\|u\|_1}{d}, \tag{5.4.14}$$

$$\operatorname{orth} \zeta_u > d, \tag{5.4.15}$$

and in particular

$$\begin{aligned} \|\zeta_u\|_\infty &\leq \max\{|\zeta_u(u)|, \|\zeta_u\|_1 - |\zeta_u(u)|\} \\ &\leq 2^d \binom{\|u\|_1}{d} \\ &\leq 2\|u\|_1{}^d. \end{aligned} \tag{5.4.16}$$

The central object of study in our proof is the following function $\zeta\colon \mathbb{N}^n \to \mathbb{R}$, built from the auxiliary objects $S(v)$ and ζ_u just introduced:

$$\zeta(x) = \sum_{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}} \Lambda(v)\, \zeta_{v|_{S(v)}}(x|_{S(v)})\, \mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}]. \tag{5.4.17}$$

The expression on the right-hand side is well-formed because, to restate (5.4.11), each string $v|_{S(v)}$ has weight greater than d and can therefore be used as a subscript in $\zeta_{v|_{S(v)}}$. Specializing (5.4.15) and (5.4.16),

$$\operatorname{orth} \zeta_{v|_{S(v)}} > d, \qquad\qquad v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}, \tag{5.4.18}$$

$$\|\zeta_{v|_{S(v)}}\|_\infty \leq 2(nr)^d, \qquad\qquad v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}. \tag{5.4.19}$$

Property (5.4.12) ensures that $\zeta_{v|_{S(v)}}(x|_{S(v)})\, \mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \neq 0$ only when $x \leq v$. It follows that

$$\begin{aligned}\operatorname{supp}\zeta &\subseteq \bigcup_{v \in \operatorname{supp}\Lambda} \{x \in \mathbb{N}^n : x \leq v\} \\ &= \operatorname{supp}\Lambda, \end{aligned} \tag{5.4.20}$$

where second step is valid because $\Lambda \in \mathfrak{B}(r, c, \alpha)^{\otimes n}$.

Before carrying on with the proof, we take a moment to simplify the defining expression for ζ. For any $v \in \mathbb{N}^n|_{\geq 2\theta}$, we have

$$\begin{aligned}
&\zeta_{v|_{S(v)}}(x|_{S(v)})\,\mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \\
&\quad = \zeta_{v|_{S(v)}}(x|_{S(v)})\,\mathbf{I}[x|_{S(v)} = v|_{S(v)} \text{ or } \|x|_{S(v)}\|_1 \leq d]\,\mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \\
&\quad = \zeta_{v|_{S(v)}}(x|_{S(v)})(\mathbf{I}[x|_{S(v)} = v|_{S(v)}] + \mathbf{I}[\|x|_{S(v)}\|_1 \leq d])\mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \\
&\quad = \zeta_{v|_{S(v)}}(x|_{S(v)})\mathbf{I}[x = v] \\
&\qquad\qquad\qquad + \zeta_{v|_{S(v)}}(x|_{S(v)})\mathbf{I}[\|x|_{S(v)}\|_1 \leq d]\,\mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \\
&\quad = \mathbf{I}[x = v] + \zeta_{v|_{S(v)}}(x|_{S(v)})\mathbf{I}[\|x|_{S(v)}\|_1 \leq d]\,\mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}],
\end{aligned}$$

where the first, second, and fourth steps are valid by (5.4.12), (5.4.11), and (5.4.13), respectively. Making this substitution in the defining equation for ζ,

$$\begin{aligned}
\zeta(x) = \sum_{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}} &\Lambda(v)\zeta_{v|_{S(v)}}(x|_{S(v)})\mathbf{I}[\|x|_{S(v)}\|_1 \leq d]\,\mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \\
&+ \sum_{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}} \Lambda(v)\mathbf{I}[x = v]. \qquad (5.4.21)
\end{aligned}$$

We proceed to establish key properties of ζ.

Step 1: Orthogonality. By Proposition 5.7(ii), each term in the summation on the right-hand side of (5.4.17) is a function orthogonal to polynomials of degree less than $\operatorname{orth}\zeta_{v|_{S(v)}}$. Therefore,

$$\begin{aligned}
\operatorname{orth}\zeta &\geq \min_{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}} \operatorname{orth}\zeta_{v|_{S(v)}} \\
&> d, \qquad (5.4.22)
\end{aligned}$$

where the first step uses Proposition 5.7(i) and the second step applies (5.4.18).

STEP 2: HEAVY INPUTS. We now examine the behavior of ζ on inputs of weight at least 2θ, which we think of as "heavy." For any string $v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}$, we have

$$\begin{aligned} x \in \mathbb{N}^n|_{\geq 2\theta} &\implies \|x\|_1 > d + \theta \\ &\implies \|x|_{S(v)}\|_1 > d \quad \vee \quad \|x|_{\overline{S(v)}}\|_1 > \theta \\ &\implies \|x|_{S(v)}\|_1 > d \quad \vee \quad x|_{\overline{S(v)}} \neq v|_{\overline{S(v)}}, \end{aligned}$$

where the final implication uses (5.4.9). We conclude that the first summation in (5.4.21) vanishes on $\mathbb{N}^n|_{\geq 2\theta}$, so that

$$\zeta(x) = \Lambda(x), \qquad\qquad x \in \mathbb{N}^n|_{\geq 2\theta}. \tag{5.4.23}$$

This completes the analysis of heavy inputs.

Step 3: Light inputs. We now turn to inputs of weight less than 2θ, the most technical part of the proof. Fix an arbitrary string $x \in (\operatorname{supp}\Lambda)|_{<2\theta}$. Then

$$\begin{aligned}
\frac{|\zeta(x)|}{\Lambda(x)} &= \left| \sum_{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}} \frac{\Lambda(v)}{\Lambda(x)} \zeta_{v|_{S(v)}}(x|_{S(v)})\, \mathbf{I}[\|x|_{S(v)}\|_1 \leq d]\, \mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \right| \\
&\leq \sum_{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}} \frac{\Lambda(v)}{\Lambda(x)} |\zeta_{v|_{S(v)}}(x|_{S(v)})|\, \mathbf{I}[\|x|_{S(v)}\|_1 \leq d]\, \mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \\
&\leq 2(nr)^d \sum_{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}} \frac{\Lambda(v)}{\Lambda(x)}\, \mathbf{I}[\|x|_{S(v)}\|_1 \leq d]\, \mathbf{I}[x|_{\overline{S(v)}} = v|_{\overline{S(v)}}] \\
&= 2(nr)^d \sum_{\substack{S \subseteq \{1,\ldots,n\}: \\ |S| \geq \theta/r}} \mathbf{I}[\|x|_S\|_1 \leq d] \sum_{\substack{v \in (\operatorname{supp}\Lambda)|_{\geq 2\theta}: \\ S(v) = S}} \frac{\Lambda(v)}{\Lambda(x)}\, \mathbf{I}[x|_{\overline{S}} = v|_{\overline{S}}] \\
&\leq 2(nr)^d \sum_{\substack{S \subseteq \{1,\ldots,n\}: \\ |S| \geq \theta/r}} \mathbf{I}[\|x|_S\|_1 \leq d] \sum_{\substack{v \in \mathbb{N}^n: \\ \sum_{i \in S} v_i \geq \theta, \\ \min_{i \in S} v_i \geq \frac{\theta}{2|S|(1+\ln n)}}} \frac{\Lambda(v)}{\Lambda(x)}\, \mathbf{I}[x|_{\overline{S}} = v|_{\overline{S}}],
\end{aligned} \tag{5.4.24}$$

where the first step uses (5.4.21); the second step applies the triangle inequality; the third step is valid by (5.4.19); the fourth step amounts to collecting terms according to $S(v)$, which by (5.4.7) has cardinality at least θ/r; and the fifth step uses (5.4.8) and (5.4.10).

Bounding (5.4.24) requires a bit of work. To start with, write $\Lambda = \bigotimes_{i=1}^n \lambda_i$ for some $\lambda_1, \lambda_2, \ldots, \lambda_n \in \mathfrak{B}(r, c, \alpha)$. Then for every nonempty set $S \subseteq \{1, 2, \ldots, n\}$,

$$\begin{aligned}
\mathbf{I}[\|x|_S\|_1 \leq d] \prod_{i \in S} \lambda_i(x_i) &\geq \mathbf{I}[\|x|_S\|_1 \leq d] \prod_{i \in S} \frac{c^{x_i+1}}{(x_i+1)^2\, 2^{\alpha x_i}} \\
&= \mathbf{I}[\|x|_S\|_1 \leq d]\, c^{|S|} \left(\frac{c}{2^\alpha}\right)^{\sum_{i\in S} x_i} \prod_{i \in S} \frac{1}{(x_i+1)^2} \\
&\geq \mathbf{I}[\|x|_S\|_1 \leq d]\, c^{|S|} \left(\frac{c}{2^\alpha}\right)^{\sum_{i\in S} x_i} \left(\frac{|S|}{\sum_{i\in S}(x_i+1)}\right)^{2|S|} \\
&\geq c^{|S|} \left(\frac{c}{2^\alpha}\right)^{d} \left(\frac{|S|}{|S|+d}\right)^{2|S|} \\
&\geq c^{|S|} \left(\frac{c}{2^\alpha e^2}\right)^{d}, \qquad (5.4.25)
\end{aligned}$$

where the first step applies the definition of $\mathfrak{B}(r, c, \alpha)$; the third step is valid by the arithmetic-geometric mean inequality; and the last step uses the bound $1 + t \leq e^t$ for real t. Continuing,

$$\begin{aligned}
&\sum_{\substack{v\in\mathbb{N}^n:\\ \sum_{i\in S} v_i\geq\theta,\\ \min_{i\in S} v_i\geq\frac{\theta}{2|S|(1+\ln n)}}} \frac{\Lambda(v)}{\Lambda(x)}\,\mathbf{I}[x|_{\overline{S}} = v|_{\overline{S}}]\\
&\qquad= \sum_{\substack{v\in\mathbb{N}^n:\\ \sum_{i\in S} v_i\geq\theta,\\ \min_{i\in S} v_i\geq\frac{\theta}{2|S|(1+\ln n)},\\ v_i=x_i \text{ for } i\notin S}} \prod_{i\in S}\frac{\lambda_i(v_i)}{\lambda_i(x_i)}\\
&\qquad\leq \sum_{\substack{v\in\mathbb{N}^n:\\ \sum_{i\in S} v_i\geq\theta,\\ \min_{i\in S} v_i\geq\frac{\theta}{2|S|(1+\ln n)},\\ v_i=x_i \text{ for } i\notin S}} 2^{-\alpha\sum_{i\in S} v_i}\prod_{i\in S}\frac{1}{c(v_i+1)^2\lambda_i(x_i)}\\
&\qquad\leq \sum_{\substack{v\in\mathbb{N}^n:\\ \min_{i\in S} v_i\geq\frac{\theta}{2|S|(1+\ln n)},\\ v_i=x_i \text{ for } i\notin S}} 2^{-\alpha\theta}\prod_{i\in S}\frac{1}{c(v_i+1)^2\lambda_i(x_i)}\\
&\qquad= 2^{-\alpha\theta}\left(\sum_{t=\lceil\frac{\theta}{2|S|(1+\ln n)}\rceil}^{\infty}\frac{1}{c(t+1)^2}\right)^{|S|}\prod_{i\in S}\frac{1}{\lambda_i(x_i)}\\
&\qquad\leq 2^{-\alpha\theta}\left(\int_{\lceil\frac{\theta}{2|S|(1+\ln n)}\rceil}^{\infty}\frac{dt}{ct^2}\right)^{|S|}\prod_{i\in S}\frac{1}{\lambda_i(x_i)}\\
&\qquad\leq 2^{-\alpha\theta}\left(\frac{2|S|(1+\ln n)}{c\theta}\right)^{|S|}\prod_{i\in S}\frac{1}{\lambda_i(x_i)}, \qquad (5.4.26)
\end{aligned}$$

where the first step uses $\Lambda = \bigotimes_{i=1}^{n}\lambda_i$, and the second step applies the definition of $\mathfrak{B}(r, c, \alpha)$.

It remains to put together the bounds obtained so far. We have:

$$\begin{aligned}
\frac{|\zeta(x)|}{\Lambda(x)} &\leq 2(nr)^d \sum_{\substack{S\subseteq\{1,\ldots,n\}:\\ |S|\geq\theta/r}} \mathbf{I}[\|x|_S\|_1 \leq d] \cdot 2^{-\alpha\theta} \left(\frac{2|S|(1+\ln n)}{c\theta}\right)^{|S|} \prod_{i\in S} \frac{1}{\lambda_i(x_i)} \\
&\leq 2(nr)^d \sum_{\substack{S\subseteq\{1,\ldots,n\}:\\ |S|\geq\theta/r}} 2^{-\alpha\theta} \left(\frac{2|S|(1+\ln n)}{c^2\theta}\right)^{|S|} \cdot \left(\frac{2^\alpha e^2}{c}\right)^d \\
&\leq 2 \cdot \frac{(e^2nr/c)^d}{2^{\alpha\lceil\theta/2\rceil}} \sum_{\substack{S\subseteq\{1,\ldots,n\}:\\ |S|\geq\theta/r}} \left(\frac{2|S|(1+\ln n)}{c^2\theta}\right)^{|S|} \\
&= 2 \cdot \frac{(e^2nr/c)^d}{2^{\alpha\lceil\theta/2\rceil}} \sum_{s=\lceil\theta/r\rceil}^{\infty} \binom{n}{s} \left(\frac{2s(1+\ln n)}{c^2\theta}\right)^{s} \\
&\leq 2 \cdot \frac{(e^2nr/c)^d}{2^{\alpha\lceil\theta/2\rceil}} \sum_{s=\lceil\theta/r\rceil}^{\infty} \left(\frac{en}{s} \cdot \frac{2s(1+\ln n)}{c^2\theta}\right)^{s} \\
&\leq 2 \cdot \frac{(e^2nr/c)^d}{2^{\alpha\lceil\theta/2\rceil}} \sum_{s=\lceil\theta/r\rceil}^{\infty} 2^{-s} \\
&= 4 \cdot \frac{(e^2nr/c)^d}{2^{\alpha\lceil\theta/2\rceil+\lceil\theta/r\rceil}},
\end{aligned}$$

where the first step follows from (5.4.24) and (5.4.26); the second step substitutes the bound from (5.4.25); the third step uses (5.4.2); and the next-to-last step uses (5.4.3). In summary, we have shown that

$$|\zeta(x)| \leq 4 \cdot \frac{(e^2nr/c)^d}{2^{\alpha\lceil\theta/2\rceil+\lceil\theta/r\rceil}}\, \Lambda(x), \qquad\qquad x \in (\operatorname{supp}\Lambda)|_{<2\theta}. \qquad (5.4.27)$$

STEP 4: FINISHING THE PROOF. Define $\tilde{\Lambda} = \Lambda - \zeta$. Then the support property (5.4.4) follows from (5.4.20) and (5.4.23); the analytic indistinguishability property (5.4.5) follows from (5.4.22); and the pointwise property (5.4.6) follows from (5.4.4) and (5.4.27). □

We record a generalization of Theorem 5.24 to translates of probability distributions in $\mathfrak{B}(r,c,\alpha)^{\otimes n}$, and further to convex combinations of such distributions. Formally, define $\mathfrak{B}(r,c,\alpha,\Delta)$ for $\Delta \geq 0$ to be the family of probability distributions λ on $\mathbb{N}$ such that $\lambda(t) \equiv \lambda'(t-a)$ for some $\lambda' \in \mathfrak{B}(r,c,\alpha)$ and $a \in [0,\Delta]$. We have:

Corollary 5.25. *Let $\Lambda \in \operatorname{conv}(\mathfrak{B}(r,c,\alpha,\Delta)^{\otimes n})$ be given, for some integers $r, \Delta \geq 0$ and reals $c > 0$ and $\alpha \geq 0$. Let d and θ be positive integers with*

$$\theta \geq 2d, \tag{5.4.28}$$

$$\theta \geq \frac{4en(1+\ln n)}{c^2}, \tag{5.4.29}$$

$$2^{\lceil \theta/r \rceil + \alpha\lceil \theta/2 \rceil} \geq 4\left(\frac{8nr}{c}\right)^d. \tag{5.4.30}$$

Then there is a probability distribution $\tilde{\Lambda}\colon \mathbb{N}^n \to \mathbb{R}$ such that

$$\operatorname{supp}\tilde{\Lambda} \subseteq (\operatorname{supp}\Lambda)|_{<2\theta+n\Delta}, \tag{5.4.31}$$

$$\operatorname{orth}(\Lambda - \tilde{\Lambda}) > d. \tag{5.4.32}$$

Proof. We first consider the special case when $\Lambda \in \mathfrak{B}(r,c,\alpha,\Delta)^{\otimes n}$. Then by definition, $\Lambda(t_1,\ldots,t_n) = \Lambda'(t_1-a_1,\ldots,t_n-a_n)$ for some probability distribution $\Lambda' \in \mathfrak{B}(r,c,\alpha)^{\otimes n}$ and integers $a_1,\ldots,a_n \in [0,\Delta]$. Applying Theorem 5.24 to Λ' yields a function $\tilde{\Lambda}'\colon \mathbb{N}^n \to \mathbb{R}$ with

$$\operatorname{supp}\tilde{\Lambda}' \subseteq (\operatorname{supp}\Lambda)|_{<2\theta}, \tag{5.4.33}$$

$$\operatorname{orth}(\Lambda' - \tilde{\Lambda}') > d, \tag{5.4.34}$$

$$|\Lambda' - \tilde{\Lambda}'| \leq \Lambda' \qquad \text{on } \operatorname{supp}\tilde{\Lambda}'. \tag{5.4.35}$$

The last property implies in particular that $\tilde{\Lambda}'$ is a nonnegative function. As a result, (5.4.32) and Proposition 5.10 guarantee that $\tilde{\Lambda}'$ is a distribution. Now the sought properties (5.4.31) and (5.4.32) follow from (5.4.33) and (5.4.34), respectively, for the probability distribution $\tilde{\Lambda}(t_1, \ldots, t_n) = \tilde{\Lambda}'(t_1 - a, \ldots, t_n - a_n)$.

In the general case of a convex combination $\Lambda = \lambda_1 \Lambda_1 + \cdots + \lambda_k \Lambda_k$ of probability distributions $\Lambda_1, \ldots, \Lambda_k \in \mathfrak{B}(r, c, \alpha, \Delta)^{\otimes n}$, one uses the technique of the previous paragraph to transform $\Lambda_1, \ldots, \Lambda_k$ individually into corresponding functions $\tilde{\Lambda}_1, \ldots, \tilde{\Lambda}_k$, and takes $\tilde{\Lambda} = \lambda_1 \tilde{\Lambda}_1 + \cdots + \lambda_k \tilde{\Lambda}_k$. □

5.4.2. A bounded dual polynomial for MP. We now turn to the construction of a gadget for our amplification theorem. Let $\mathfrak{B}^*(r, c, \alpha)$ denote the family of probability distributions λ on $\mathbb{N}$ such that

$$\operatorname{supp} \lambda = \{0, 1, 2, \ldots, r'\}$$

for some nonnegative integer $r' \leq r$, and moreover

$$\frac{c}{(t+1)^2 \, 2^{\alpha t}} \leq \lambda(t) \leq \frac{1}{c(t+1)^2 \, 2^{\alpha t}}, \qquad\qquad t \in \operatorname{supp} \lambda.$$

In this family, a distribution's weight at any given point is prescribed up to the multiplicative constant c, in contrast to the exponentially large range allowed in the definition of $\mathfrak{B}(r, c, \alpha)$. For all parameter settings, we have

$$\mathfrak{B}^*(r, c, \alpha) \subseteq \mathfrak{B}(r, c, \alpha). \tag{5.4.36}$$

Indeed, the containment holds trivially for $c \leq 1$, and remains valid for $c > 1$ because the left-hand side and right-hand side are both empty in that case. As before, it will be helpful to have shorthand notation for *translates* of distributions in $\mathfrak{B}(r, c, \alpha)$: we define $\mathfrak{B}^*(r, c, \alpha, \Delta)$ for $\Delta \geq 0$ to be the family of probability distributions λ on $\mathbb{N}$ such that $\lambda(t) = \lambda'(t - a)$ for some $\lambda' \in \mathfrak{B}^*(r, c, \alpha)$ and $a \in [0, \Delta]$.

As a first step toward analyzing the threshold degree of $\mathbf{AC}^0$, we will construct a dual object that witnesses the high threshold degree of $\mathrm{MP}^*_{m,r}$ and possesses additional metric properties in the sense of $\mathfrak{B}^*$. To simplify the exposition, we start with an auxiliary construction.

LEMMA 5.26. *Let $0 < \varepsilon < 1$ be given. Then for some constants $c_1, c_2 \in (0,1)$ and all integers $R \geq r \geq 1$, there are (explicitly given) probability distributions $\lambda_0, \lambda_1, \lambda_2$ such that:*

$$\operatorname{supp}\lambda_0 = \{0\}, \tag{5.4.37}$$

$$\operatorname{supp}\lambda_i = \{1,2,\ldots,R\}, \qquad i = 1,2, \tag{5.4.38}$$

$$\lambda_i \in \mathfrak{B}^*\left(R, c_1, \frac{c_2}{\sqrt{r}}, 1\right), \qquad i = 0,1,2, \tag{5.4.39}$$

$$\operatorname{orth}((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2 - \lambda_1) \geq c_1\sqrt{r}. \tag{5.4.40}$$

Our analysis of the threshold degree of $\mathbf{AC}^0$ only uses the special case $R = r$ of Lemma 5.26. The more general formulation with $R \geq r$ will be needed much later, in the analysis of the sign-rank of $\mathbf{AC}^0$.

Proof. Theorem 5.17 constructs a function $\psi\colon \{0,1,2,\ldots,R\} \to \mathbb{R}$ such that

$$\psi(0) > \frac{1-\frac{\varepsilon}{2}}{2}, \tag{5.4.41}$$

$$\|\psi\|_1 = 1, \tag{5.4.42}$$

$$\operatorname{orth}\psi \geq c'\sqrt{r}, \tag{5.4.43}$$

$$|\psi(t)| \in \left[\frac{c'}{(t+1)^2\, 2^{c''t/\sqrt{r}}}, \frac{1}{c'(t+1)^2\, 2^{c''t/\sqrt{r}}}\right], \qquad t = 0,1,\ldots,r, \tag{5.4.44}$$

for some absolute constants $c', c'' \in (0,1)$. Property (5.4.42) makes it possible to view $|\psi|$ as a probability distribution on $\{0,1,2,\ldots,R\}$. Let μ_0, μ_1, μ_2 be the probability distributions induced by $|\psi|$ on $\{0\}, \{t \neq 0 : \psi(t) < 0\}$, and $\{t \neq 0 : \psi(t) > 0\}$, respectively. It is clear from (5.4.41) that the negative part of ψ is a multiple of μ_1, whereas the positive part of ψ is a nonnegative linear combination of μ_0 and μ_2. Moreover, it follows from $\langle \psi, 1 \rangle = 0$ and $\|\psi\|_1 = 1$ that the positive and negative parts of ψ both have ℓ_1-norm 1/2. Summarizing,

$$\psi = \frac{1-\delta}{2}\mu_0 - \frac{1}{2}\mu_1 + \frac{\delta}{2}\mu_2 \tag{5.4.45}$$

for some $0 \leq \delta \leq 1$. In view of (5.4.41), we infer the more precise bound

$$0 \leq \delta < \frac{\varepsilon}{2}. \tag{5.4.46}$$

We define

$$\lambda_0 = \mu_0, \tag{5.4.47}$$

$$\lambda_1 = \frac{1-\varepsilon\delta}{1-\delta^2}\mu_1 + \delta \cdot \frac{\varepsilon - \delta}{1-\delta^2}\mu_2, \tag{5.4.48}$$

$$\lambda_2 = \frac{\varepsilon-\delta}{\varepsilon(1-\delta^2)}\mu_1 + \delta \cdot \frac{1-\varepsilon\delta}{\varepsilon(1-\delta^2)}\mu_2. \tag{5.4.49}$$

It follows from $0 \leq \delta \leq \varepsilon$ that λ_1 and λ_2 are convex combinations of μ_1 and μ_2 and are therefore probability distributions with support

$$\operatorname{supp} \lambda_i \subseteq \{1, 2, \ldots, R\}, \qquad i = 1, 2. \tag{5.4.50}$$

Recall from (5.4.45) that $|\psi| = \frac{1}{2}\mu_1 + \frac{\delta}{2}\mu_2$ on $\{1, 2, \ldots, R\}$. Comparing the coefficients in $|\psi| = \frac{1}{2}\mu_1 + \frac{\delta}{2}\mu_2$ with the corresponding coefficients in the defining equations for λ_1 and λ_2, where $0 \leq \delta \leq \varepsilon/2$ by (5.4.46), we conclude that $\lambda_1, \lambda_2 \in [c'''|\psi|, |\psi|/c''']$ on $\{1, 2, \ldots, R\}$ for some constant $c''' = c'''(\varepsilon) \in (0, 1)$. In view of (5.4.44), we arrive at

$$|\lambda_i(t)| \in \left[\frac{c'c'''}{(t+1)^2\,2^{c''t/\sqrt{r}}},\ \frac{1}{c'c'''(t+1)^2\,2^{c''t/\sqrt{r}}}\right],$$

$$i = 1, 2; \quad t = 1, 2, \ldots, R. \tag{5.4.51}$$

Continuing,

$$\begin{aligned}\operatorname{orth}((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2 - \lambda_1) &= \operatorname{orth}\left(2 \cdot \frac{1-\varepsilon}{1-\delta}\left(\frac{1-\delta}{2}\mu_0 - \frac{1}{2}\mu_1 + \frac{\delta}{2}\mu_2\right)\right) \\ &= \operatorname{orth}\left(2 \cdot \frac{1-\varepsilon}{1-\delta}\,\psi\right) \\ &\geq c'\sqrt{r}, \end{aligned} \tag{5.4.52}$$

where the first step follows from the defining equations (5.4.47)–(5.4.49), the second step uses (5.4.45), and the final step is a restatement of (5.4.46).

We are now in a position to verify the claimed properties of $\lambda_0, \lambda_1, \lambda_2$ in the theorem statement. Property (5.4.37) follows from (5.4.47), whereas property (5.4.38) is immediate from (5.4.50) and (5.4.51). The remaining properties (5.4.39) and (5.4.40) for small enough constants $c_1, c_2 \in (0, 1)$ now follow from (5.4.51) and (5.4.52), respectively. □

We are now in a position to construct our desired dual polynomial for the Minsky–Papert function.

Theorem 5.27. *For some absolute constants $c_1, c_2 \in (0,1)$ and all positive integers m and r, there are probability distributions Λ_0, Λ_1 such that*

$$\Lambda_i \in \operatorname{conv}\left(\mathfrak{B}^*\left(r, c_1, \frac{c_2}{\sqrt{r}}, 1\right)^{\otimes m}\right), \qquad i = 0,1, \tag{5.4.53}$$

$$\operatorname{supp}\Lambda_i \subseteq (\mathrm{MP}^*_{m,r})^{-1}(i), \qquad i = 0,1, \tag{5.4.54}$$

$$\operatorname{orth}(\Lambda_1 - \Lambda_0) \geq \min\{m, c_1\sqrt{r}\}. \tag{5.4.55}$$

The last two properties in the theorem statement are equivalent, in the sense of linear programming duality, to the lower bound $\deg_\pm(\mathrm{MP}^*_{m,r}) \geq \min\{m, c_1\sqrt{r}\}$ and can be recovered in a black-box manner from many previous papers, e.g., [**90, 114, 122**]. The key new property that we prove is (5.4.53), with the newly established Lemma 5.26 playing an essential role.

Proof of Theorem 5.27. Take $\varepsilon = 1/2$ and $R = r$ in Lemma 5.26, and let $\lambda_0, \lambda_1, \lambda_2$ be the resulting probability distributions. Let

$$\Lambda_0 = \mathop{\mathbf{E}}_{\substack{S\subseteq\{1,2,\ldots,m\}\\ |S| \text{ odd}}} \lambda_0^{\otimes S} \cdot \lambda_2^{\otimes \overline{S}},$$

$$\Lambda_1 = \lambda_1^{\otimes m}.$$

Then (5.4.53) is immediate from (5.4.39), whereas (5.4.54) follows from (5.4.37) and (5.4.38). To verify the remaining property (5.4.55), rewrite

$$\begin{aligned}\Lambda_0 &= 2^{-m+1} \sum_{\substack{S\subseteq\{1,2,\ldots,m\}\\ |S| \text{ odd}}} \lambda_0^{\otimes S} \cdot \lambda_2^{\otimes \overline{S}} \\ &= \left(\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)^{\otimes m} - \left(-\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)^{\otimes m}.\end{aligned}$$

Observe that

$$\operatorname{orth}(\lambda_i - \lambda_j) \geq 1, \qquad\qquad i, j = 0, 1, 2, \tag{5.4.56}$$

which can be seen from $\langle \lambda_i - \lambda_j, 1\rangle = \langle \lambda_i, 1\rangle - \langle \lambda_j, 1\rangle = 1 - 1 = 0$. Now

$$\begin{aligned}
&\operatorname{orth}(\Lambda_1 - \Lambda_0)\\
&\quad= \operatorname{orth}\left(\lambda_1^{\otimes m} - \left(\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)^{\otimes m} + \left(-\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)^{\otimes m}\right)\\
&\quad\geq \min\left\{\operatorname{orth}\left(\lambda_1^{\otimes m} - \left(\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)^{\otimes m}\right), \operatorname{orth}\left(-\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)^{\otimes m}\right\}\\
&\quad\geq \min\left\{\operatorname{orth}\left(\lambda_1 - \frac{1}{2}\lambda_0 - \frac{1}{2}\lambda_2\right), \operatorname{orth}\left(-\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)^{\otimes m}\right\}\\
&\quad= \min\left\{\operatorname{orth}\left(\lambda_1 - \frac{1}{2}\lambda_0 - \frac{1}{2}\lambda_2\right), m \operatorname{orth}\left(-\frac{1}{2}\lambda_0 + \frac{1}{2}\lambda_2\right)\right\}\\
&\quad= \min\left\{\operatorname{orth}\left(\lambda_1 - \frac{1}{2}\lambda_0 - \frac{1}{2}\lambda_2\right), m\right\}\\
&\quad\geq \min\{c\sqrt{r}, m\},
\end{aligned}$$

where the last four steps are valid by Proposition 5.7(i), Proposition 5.7(iii), Proposition 5.7(ii), equation (5.4.56), and equation (5.4.40), respectively. □

5.4.3. Hardness amplification for threshold degree and beyond. We now present a blackbox transformation that takes any given circuit with threshold degree $n^{1-\varepsilon}$ into a circuit with polynomially larger threshold degree, $\Omega(n^{1-\frac{\varepsilon}{1+\varepsilon}})$. This hardness amplification procedure increases the circuit size additively by $n^{O(1)}$ and the circuit depth by 2, preserving membership in $\mathbf{AC}^0$. We obtain analogous hardness amplification results for a host of other approximation-theoretic complexity measures. For this reason, we adopt the following abstract view of polynomial approximation. Let I_0, I_1, I_* be nonempty convex subsets of the real line, i.e., any kind of nonempty intervals (closed, open, or half-open; bounded or unbounded). Let $f\colon X \to \{0,1,*\}$ be a (possibly partial) Boolean function on a finite subset X of Euclidean space. We define an *(I_0, I_1, I_*)-approximant for f* to be any real polynomial p that maps $f^{-1}(0), f^{-1}(1), f^{-1}(*)$ into I_0, I_1, I_*, respectively. The *(I_0, I_1, I_*)-approximate degree of f*, denoted $\deg_{I_0,I_1,I_*}(f)$, the least degree of an (I_0, I_1, I_*)-approximant for f. Threshold degree corresponds to the special case

$$\deg_{\pm} = \deg_{(0,\infty),(-\infty,0),(-\infty,\infty)} . \tag{5.4.57}$$

Other notable cases include *ε-approximate degree* and *one-sided ε-approximate degree*, given by

$$\deg_{\varepsilon} = \deg_{[-\varepsilon,\varepsilon],[1-\varepsilon,1+\varepsilon],[-\varepsilon,1+\varepsilon]}, \tag{5.4.58}$$

$$\deg_{\varepsilon}^{+} = \deg_{[-\varepsilon,\varepsilon],[1-\varepsilon,\infty),(-\infty,\infty)}, \tag{5.4.59}$$

respectively. Our hardness amplification result applies to (I_0, I_1, I_*)-approximate degree for any nonempty convex $I_0, I_1, I_* \subseteq \mathbb{R}$, with threshold degree being a special case. The centerpiece of our argument is the following lemma.

LEMMA 5.28. *Let $c, c', c'' > 0$ be the absolute constants from Theorem 5.27. Let n, m, r, d, θ be positive integers such that*

$$\theta \geq 2d, \tag{5.4.60}$$

$$\theta \geq \frac{4enm(1+\ln(nm))}{c'^2}, \tag{5.4.61}$$

$$\theta \geq \frac{2\sqrt{r}}{c''}\left(d\log\left(\frac{8nmr}{c'}\right)+2\right). \tag{5.4.62}$$

Then for each $z \in \{0,1\}^n$, there is a probability distribution $\tilde{\Lambda}_z$ on $\mathbb{N}^{nm}$ such that:

(i) *the support of $\tilde{\Lambda}_z$ is contained in $(\prod_{i=1}^{n}(\mathrm{MP}^{*}_{m,r})^{-1}(z_i))|_{<2\theta+nm}$;*

(ii) *for every polynomial* $p\colon \mathbb{R}^{nm} \to \mathbb{R}$ *of degree at most* d, *the mapping* $z \mapsto \mathbf{E}_{\tilde{\Lambda}_z} p$ *is a polynomial on* $\{0,1\}^n$ *of degree at most* $\frac{1}{\min\{m, c\sqrt{r}\}} \cdot \deg p$.

Proof. Theorem 5.27 constructs probability distributions Λ_0 and Λ_1 such that

$$\Lambda_i \in \operatorname{conv}\left(\mathfrak{B}^*\left(r, c', \frac{c''}{\sqrt{r}}, 1\right)^{\otimes m}\right), \qquad i = 0, 1, \tag{5.4.63}$$

$$\operatorname{supp} \Lambda_i \subseteq (\mathrm{MP}^*_{m,r})^{-1}(i), \qquad i = 0, 1, \tag{5.4.64}$$

$$\operatorname{orth}(\Lambda_1 - \Lambda_0) \geq \min\{m, c\sqrt{r}\}. \tag{5.4.65}$$

As a result, the probability distributions $\Lambda_z = \bigotimes_{i=1}^n \Lambda_{z_i}$ for $z \in \{0,1\}^n$ obey

$$\begin{aligned}\Lambda_z &\in \left(\operatorname{conv}\left(\mathfrak{B}^*\left(r, c', \frac{c''}{\sqrt{r}}, 1\right)^{\otimes m}\right)\right)^{\otimes n} \\ &\subseteq \operatorname{conv}\left(\mathfrak{B}^*\left(r, c', \frac{c''}{\sqrt{r}}, 1\right)^{\otimes nm}\right) \\ &\subseteq \operatorname{conv}\left(\mathfrak{B}\left(r, c', \frac{c''}{\sqrt{r}}, 1\right)^{\otimes nm}\right).\end{aligned} \tag{5.4.66}$$

By (5.4.60)–(5.4.62), (5.4.66), and Corollary 5.25, there are probability distribution $\tilde{\Lambda}_z\colon \mathbb{N}^{nm} \to \mathbb{R}$ for $z \in \{0,1\}^n$ such that

$$\operatorname{supp} \tilde{\Lambda}_z \subseteq (\operatorname{supp} \Lambda_z)|_{<2\theta+nm}, \tag{5.4.67}$$

$$\operatorname{orth}(\Lambda_z - \tilde{\Lambda}_z) > d, \tag{5.4.68}$$

We proceed to verify the properties required of $\tilde{\Lambda}_z$. For (i), it follows from (5.4.64) and (5.4.67) that each $\tilde{\Lambda}_z$ has support contained in $(\prod_{i=1}^n (\mathrm{MP}^*_{m,r})^{-1}(z_i))|_{<2\theta+nm}$. For (ii), let p be any polynomial of degree at most d. Then (5.4.68) guarantees that $\mathbf{E}_{\tilde{\Lambda}_z} p = \mathbf{E}_{\Lambda_z} p$, where the right-hand side is by (5.4.65) and Proposition 5.8 a polynomial in $z \in \{0,1\}^n$ of degree at most $\deg p / \operatorname{orth}(\Lambda_1 - \Lambda_0) \leq \deg p / \min\{m, c\sqrt{r}\}$. □

At its core, a hardness amplification result is a lower bound on the complexity of a composed function in terms of the complexities of its constituent parts. We now prove such a composition theorem for (I_0, I_1, I_*)-approximate degree.

THEOREM 5.29. *There is an absolute constant* $0 < c < 1$ *such that*

$$\deg_{I_0,I_1,I_*}((f \circ \mathrm{MP}^*_m)|_{\leq\theta}) \geq \min\left\{cm\deg_{I_0,I_1,I_*}(f), \frac{c\theta}{m\log(n+m)} - n\right\},$$

$$\deg_{I_0,I_1,I_*}((f \circ \neg\mathrm{MP}^*_m)|_{\leq\theta}) \geq \min\left\{cm\deg_{I_0,I_1,I_*}(f), \frac{c\theta}{m\log(n+m)} - n\right\}$$

for all positive integers $n, m, \theta,$ *all functions* $f\colon \{0,1\}^n \to \{0,1,*\}$, *and all nonempty convex sets* $I_0, I_1, I_* \subseteq \mathbb{R}$.

As a practical matter, note that the left-hand sides of the inequalities in Theorem 5.29 are monotonic functions of m. Therefore, the theorem implies that $(f \circ \mathrm{MP}^*_m)|_{\leq\theta}$ and $(f \circ \neg\mathrm{MP}^*_m)|_{\leq\theta}$ have (I_0, I_1, I_*)-approximate degree at least

$$\max_{m'=1,2,\ldots,m} \min\left\{cm'\deg_{I_0,I_1,I_*}(f), \frac{c\theta}{m'\log(n+m')} - n\right\}.$$

Proof of Theorem 5.29. Negating a function's input has no effect on the (I_0, I_1, I_*)-approximate degree, so that $f(x_1, x_2, \ldots, x_n)$ and $f(\neg x_1, \neg x_2, \ldots, \neg x_n)$ both have (I_0, I_1, I_*)-approximate degree $\deg_{I_0,I_1,I_*}(f)$. Therefore, it suffices to prove the lower bound on $\deg_{I_0,I_1,I_*}((f \circ \mathrm{MP}^*_m)|_{\leq\theta})$ for all f.

Let $c \in (0,1)$ be an absolute constant that is sufficiently small relative to the constants in Lemma 5.28. For $\theta \leq \frac{1}{c} \cdot nm\log(n+m)$, the lower bounds in the statement of the theorem are nonpositive and therefore trivially true. In the complementary case $\theta > \frac{1}{c} \cdot nm\log(n+m)$, Lemma 5.28 applies to the positive integers n', m', r', d', θ', where

$$\begin{aligned} n' &= n, \\ m' &= m, \\ r' &= m^2, \\ \theta' &= \left\lfloor \frac{\theta - nm}{2} \right\rfloor, \\ d' &= \left\lfloor \frac{c\theta}{m\log(n+m)} \right\rfloor. \end{aligned}$$

We thus obtain, for each $z \in \{0,1\}^n$, a probability distribution $\tilde{\Lambda}_z$ on $\mathbb{N}^{nm}$ such that: □

(i) the support of $\tilde{\Lambda}_z$ is contained in $(\prod_{i=1}^n (\mathrm{MP}_m^*)^{-1}(z_i))|_{\leq\theta}$;

(ii) for every polynomial $p\colon \mathbb{R}^{nm} \to \mathbb{R}$ of degree at most d', the mapping $z \mapsto \mathbf{E}_{\tilde{\Lambda}_z}\, p$ is a polynomial on $\{0,1\}^n$ of degree at most $\frac{1}{cm} \cdot \deg p$.

Proof. Now, let $p\colon \mathbb{R}^{nm} \to \mathbb{R}$ be an (I_0, I_1, I_*)-approximant for $(f \circ \mathrm{MP}_m^*)|_{\leq\theta}$ of degree at most d'. Consider the mapping $p^*\colon z \mapsto \mathbf{E}_{\tilde{\Lambda}_z}\, p$, which we view as a polynomial in $z \in \{0,1\}^n$. Then (i) along with the convexity of I_0, I_1, I_* ensures that p^* is an (I_0, I_1, I_*)-approximant for f, whence $\deg p^* \geq \deg_{I_0,I_1,I_*}(f)$. At the same time, (ii) guarantees that $\deg p^* \leq \frac{1}{cm} \cdot \deg p$. This pair of lower and upper bounds force

$$\deg p \geq cm \deg_{I_0,I_1,I_*}(f).$$

Since p was chosen arbitrarily from among (I_0, I_1, I_*)-approximants of $(f \circ \mathrm{MP}_m^*)|_{\leq\theta}$ that have degree at most d', we conclude that

$$\begin{aligned} \deg_{I_0,I_1,I_*}((f \circ \mathrm{MP}_m)|_{\leq\theta}) &\geq \min\{cm \deg_{I_0,I_1,I_*}(f),\ d'+1\} \\ &\geq \min\left\{cm \deg_{I_0,I_1,I_*}(f),\ \frac{c\theta}{m\log(n+m)}\right\}. \end{aligned}$$ □

The previous composition theorem has the following analogue for Boolean inputs.

THEOREM 5.30. *Let $0 < c < 1$ be the absolute constant from Theorem 5.29. Let n, m, N be positive integers. Then there is an (explicitly given) transformation $H\colon \{0,1\}^N \to \{0,1\}^n$, computable by an AND-OR-AND circuit of size $(Nnm)^{O(1)}$ with bottom fan-in $O(\log(nm))$, such that for all functions $f\colon \{0,1\}^n \to \{0,1,*\}$ and all nonempty convex sets $I_0, I_1, I_* \subseteq \mathbb{R}$,*

$$\deg_{I_0,I_1,I_*}(f \circ H) \geq \min\left\{cm \deg_{I_0,I_1,I_*}(f), \frac{cN}{50m\log^2(n+m)} - n\right\} \log(n+m),$$

$$\deg_{I_0,I_1,I_*}(f \circ \neg H) \geq \min\left\{cm \deg_{I_0,I_1,I_*}(f), \frac{cN}{50m\log^2(n+m)} - n\right\} \log(n+m).$$

Proof. As in the previous proof, settling the first lower bound for all f will automatically settle the second lower bound, due to the invariance of (I_0, I_1, I_*)-approximate degree under negation of the input bits. In what follows, we focus on $f \circ H$.

We may assume that $N \geq 50mn\log^2(n+m)$ since otherwise the lower bounds in the theorem statement are nonpositive and hence trivially true. Define

$$\theta = \left\lceil \frac{N}{50\log(n+m)} \right\rceil.$$

Theorem 5.23 gives a surjection $G\colon \{0,1\}^{6\theta\lceil\log(nm+1)\rceil} \to \mathbb{N}^{nm}|_{\leq\theta}$ with the following two properties:

(i) for every coordinate $i = 1, 2, \ldots, nm$, the mapping $x \mapsto \mathrm{OR}^*_\theta(G(x)_i)$ is computable by an explicit DNF formula of size $(nm\theta)^{O(1)} = N^{O(1)}$ with bottom fan-in $O(\log(nm))$;

(ii) for any polynomial p, the map $v \mapsto \mathbf{E}_{G^{-1}(v)}\, p$ is a polynomial on $\mathbb{N}^{nm}|_{\leq\theta}$ of degree at most $(\deg p)/\lceil\log(nm+1)+1\rceil \leq (\deg p)/\log(n+m)$.

Consider the composition $F = (f \circ \mathrm{MP}^*_{m,\theta}) \circ G$. Then

$$\begin{aligned} F &= (f \circ (\mathrm{AND}_m \circ \mathrm{OR}^*_\theta)) \circ G \\ &= f \circ ((\underbrace{\mathrm{AND}_m \circ \mathrm{OR}^*_\theta, \ldots, \mathrm{AND}_m \circ \mathrm{OR}^*_\theta}_{n}) \circ G), \end{aligned}$$

which by property (i) of G means that F is the composition of f and an AND-OR-AND circuit H on $6\theta\lceil\log(nm+1)\rceil \le N$ variables of size $(nmN)^{O(1)} = N^{O(1)}$ with bottom fan-in $O(\log(nm))$. Hence, the proof will be complete once we show that

$$\deg_{I_0,I_1,I_*}(F) \ge \min\left\{cm\deg_{I_0,I_1,I_*}(f), \frac{cN}{50m\log^2(n+m)} - n\right\}\log(n+m). \tag{5.4.69}$$

For this, fix an (I_0, I_1, I_*)-approximant p for F of degree $\deg_{I_0,I_1,I_*}(F)$. Consider the polynomial $p^*\colon \mathbb{N}^{nm}|_{\le\theta} \to \mathbb{R}$ given by $p^*(v) = \mathbf{E}_{G^{-1}(v)}\, p$. Since I_0, I_1, I_* are convex and p is an (I_0, I_1, I_*)-approximant for $F = (f \circ \mathrm{MP}^*_{m,\theta}) \circ G$, it follows that p^* is an (I_0, I_1, I_*)-approximant for $(f \circ \mathrm{MP}^*_{m,\theta})|_{\le\theta}$. Therefore,

$$\begin{aligned} \deg p^* &\ge \deg_{I_0,I_1,I_*}((f \circ \mathrm{MP}^*_{m,\theta})|_{\le\theta}) \\ &\ge \deg_{I_0,I_1,I_*}((f \circ \mathrm{MP}^*_{m})|_{\le\theta}) \\ &\ge \min\left\{cm\deg_{I_0,I_1,I_*}(f), \frac{c\theta}{m\log(n+m)} - n\right\} \\ &\ge \min\left\{cm\deg_{I_0,I_1,I_*}(f), \frac{cN}{50m\log^2(n+m)} - n\right\}, \end{aligned}$$

where the second step is valid because $\mathrm{MP}^*_{m,\theta}$ contains $\mathrm{MP}^*_m = \mathrm{MP}^*_{m,m^2}$ as a subfunction, and the third step is legitimate by Theorem 5.29. However, property (ii) of G

states that

$$\begin{aligned}\deg p^* &\leq \frac{\deg p}{\log(n+m)}\\ &= \frac{\deg_{I_0,I_1,I_*}(F)}{\log(n+m)}.\end{aligned}$$

Comparing these lower and upper bounds on the degree of p^* settles (5.4.69). □

At last, we illustrate the use of the previous two composition results to amplify hardness for polynomial approximation.

THEOREM 5.31 (Hardness amplification). *Let $I_0, I_1, I_* \subseteq \mathbb{R}$ be any nonempty convex subsets. Let $f\colon \{0,1\}^n \to \{0,1\}$ be a given function with*

$$\deg_{I_0,I_1,I_*}(f) \geq n^{1-\frac{1}{k}},$$

for some real number $k \geq 1$. Suppose further that f is computable by a Boolean circuit of size s and depth d, where $d \geq 1$. Then there is a function $F\colon \{0,1\}^N \to \{0,1\}$ on $N = \Theta(n^{1+\frac{1}{k}} \log^2 n)$ variables with

$$\deg_{I_0,I_1,I_*}(F) \geq \Omega\left(\frac{N^{1-\frac{1}{k+1}}}{\log^{1-\frac{2}{k+1}} N}\right).$$

Moreover, F is computable by a Boolean circuit of size $s + n^{O(1)}$, bottom fan-in $O(\log n)$, depth $d+2$ if the circuit for f is monotone, and depth $d+3$ otherwise.

Proof. Take

$$\begin{aligned}m &= \lceil n^{1/k} \rceil,\\ N &= \left\lceil \frac{100}{c} mn \log^2(n+m) \right\rceil,\end{aligned}$$

where $0 < c < 1$ is the absolute constant from Theorem 5.29. Then Theorem 5.30 gives an explicit transformation $H\colon \{0,1\}^N \to \{0,1\}^n$, computable by an AND-OR-AND circuit of size $n^{O(1)}$ with bottom fan-in $O(\log n)$, such that

$$\begin{aligned}
&\min\{\deg_{I_0,I_1,I_*}(f \circ H), \deg_{I_0,I_1,I_*}(f \circ \neg H)\} \\
&\qquad \geq \min\left\{cm \deg_{I_0,I_1,I_*}(f), \frac{cN}{50m\log^2(n+m)} - n\right\} \log(n+m) \\
&\qquad \geq cn\log n \\
&\qquad = \Theta\left(\frac{N^{1-\frac{1}{k+1}}}{\log^{1-\frac{2}{k+1}} N}\right).
\end{aligned}$$

Now, fix a circuit for f of size s and depth $d \geq 1$. Composing the circuits for f and H results in circuits for $f \circ H$ and $f \circ \neg H$ of size $s + n^{O(1)}$, bottom fan-in $O(\log n)$, and depth at most $d+3$. Thus, F can be taken to be either of $f \circ H$ and $f \circ \neg H$. When the circuit for f is monotone, the depth of F can be reduced to $d+2$ as follows. After merging like gates if necessary, the circuit for f can be viewed as composed of d layers of alternating gates ($\wedge$ and $\vee$). The bottom layer of f can therefore be merged with the top layer of either H or $\neg H$, resulting in a circuit of depth at most $d+3-1=2$. □

We emphasize that in view of (5.4.57), the symbol $\deg_{I_0,I_1,I_*}$ in Theorems 5.29–5.31 can be replaced with the threshold degree symbol $\deg_\pm$. The same goes for any other special case of (I_0, I_1, I_*)-approximate degree.

5.4.4. Threshold degree of surjectivity. We start with the simplest application of our amplification theorem, in which the outer function f is the identity map $f\colon \{0,1\} \to \{0,1\}$ on a single bit.

THEOREM 5.32. *For any integer $m \geq 1$,*

$$\deg_{\pm}(\mathrm{MP}^*_m|_{\leq m^2 \log m}) = \Omega(m).$$

Proof. Let $f\colon \{0,1\} \to \{0,1\}$ be the identity function, so that $\deg_{\pm}(f) = 1$. Invoking Theorem 5.29 with $n = 1$ and $\theta = \lfloor m^2 \log m \rfloor$, one obtains the claimed lower bound. □

Theorem 5.32 has a useful interpretation. For positive integers n and r, the *surjectivity problem* is the problem of determining whether a given mapping $\{1, 2, \ldots, n\} \to \{1, 2, \ldots, r\}$ is surjective. This problem is trivial for $r > n$, and the standard regime studied in previous work is $r \leq cn$ for some constant $0 < c < 1$. The input to the surjectivity problem is represented by a Boolean matrix $x \in \{0,1\}^{r \times n}$ with precisely one nonzero entry in every column. More formally, let $e_1, e_2, \ldots, e_r$ be the standard basis for $\mathbb{R}^n$. The surjectivity function $\mathrm{SURJ}_{n,r}\colon \{e_1, e_2, \ldots, e_r\}^n \to \{0,1\}$ is given by

$$\mathrm{SURJ}_{n,r}(x_1, x_2, \ldots, x_n) = \bigwedge_{j=1}^{r} \bigvee_{i=1}^{n} x_{i,j}.$$

It is clear that $\mathrm{SURJ}_{n,r}(x_1, x_2, \ldots, x_n)$ is uniquely determined by the vector sum $x_1 + x_2 + \cdots + x_n \in \mathbb{N}^r|_n$. It is therefore natural to consider a symmetric counterpart of the surjectivity function, with domain $\mathbb{N}^r|_n$ instead of $\{e_1, e_2, \ldots, e_r\}^n$. This symmetric version is $(\mathrm{AND}_r \circ \mathrm{OR}^*_n)|_n = \mathrm{MP}^*_{r,n}|_n$, and Proposition 5.13 ensures that

$$\deg_{\pm}(\mathrm{SURJ}_{n,r}) = \deg_{\pm}(\mathrm{MP}^*_{r,n}|_n). \tag{5.4.70}$$

The surjectivity problem has seen much work recently [**15, 126, 33, 39**]. In particular, Bun and Thaler [**39**] have obtained an essentially tight lower bound of $\tilde{\Omega}(\min\{r, \sqrt{n/\log n}\})$ on the threshold degree of $\mathrm{SURJ}_{n,r}$ in the standard regime $r \leq (1 - \Omega(1))n$. As a corollary to Theorem 5.32, we give a new proof of Bun and Thaler's result, sharpening their bound by a polylogarithmic factor.

COROLLARY 5.33. *For any integers $n > r \geq 1$,*

$$\deg_\pm(\mathrm{SURJ}_{n,r}) \geq \Omega\left(\min\left\{r, \sqrt{\frac{n-r}{1+\log(n-r)}}\right\}\right). \tag{5.4.71}$$

Proof. Define

$$r' = \min\left\{r-1, \left\lfloor\sqrt{\frac{n-r}{1+\log(n-r)}}\right\rfloor\right\}. \tag{5.4.72}$$

We may assume that $r' \geq 1$ since (5.4.71) holds trivially otherwise. The identity

$$\begin{aligned}&\mathrm{MP}^*_{r',n}(x_1, x_2, \ldots, x_{r'})\\ &\qquad\qquad = \mathrm{MP}^*_{r,n}\left(x_1, x_2, \ldots, x_{r'}, \underbrace{1, 1, \ldots, 1}_{r-r'-1}, 1 + n - (r - r') - \sum_{i=1}^{r'} x_i\right)\end{aligned}$$

holds for all $(x_1, x_2, \ldots, x_{r'}) \in \mathbb{N}^{r'}|_{\leq n-(r-r')}$, whence

$$\deg_\pm(\mathrm{MP}^*_{r',n}|_{\leq n-(r-r')}) \leq \deg_\pm(\mathrm{MP}^*_{r,n}|_n). \tag{5.4.73}$$

Now

$$\begin{aligned}\deg_\pm(\mathrm{SURJ}_{n,r}) &= \deg_\pm(\mathrm{MP}^*_{r,n}|_n)\\ &\geq \deg_\pm(\mathrm{MP}^*_{r',n}|_{\leq n-(r-r')})\\ &\geq \deg_\pm(\mathrm{MP}^*_{r',r'^2}|_{\leq r'^2 \log r'})\\ &\geq \Omega(r'),\end{aligned}$$

where the four steps use (5.4.70), (5.4.73), (5.4.72), and Theorem 5.32, respectively.

□

5.4.5. Threshold degree and discrepancy of $\mathbf{AC^0}$.

We now turn to our main result on the sign-representation of constant-depth circuits. For any $\varepsilon > 0$, the next theorem constructs a circuit family in $\mathbf{AC}^0$ with threshold degree $\Omega(n^{1-\varepsilon})$. The proof amounts to a recursive application of the hardness amplification procedure of Section 5.4.3.

THEOREM 5.34. *Let $k \geq 1$ be a fixed integer. Then there is an (explicitly given) family of functions $\{f_{k,n}\}_{n=1}^{\infty}$, where $f_{k,n}\colon \{0,1\}^n \to \{0,1\}$ has threshold degree*

$$\deg_{\pm}(f_{k,n}) = \Omega\left(n^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right) \tag{5.4.74}$$

and is computable by a monotone Boolean circuit of size $n^{O(1)}$ and depth k. In addition, the circuit for $f_{k,n}$ has bottom fan-in $O(\log n)$ for all $k \neq 2$.

Proof. The proof is by induction on k. The base cases $k = 1$ and $k = 2$ correspond to the families

$$f_{1,n}(x) = x_1, \qquad\qquad n = 1,2,3,\ldots,$$
$$f_{2,n}(x) = \mathrm{MP}_{\lfloor n^{1/3}\rfloor}, \qquad\qquad n = 1,2,3,\ldots.$$

For the former, the threshold degree lower bound (5.4.74) is trivial. For the latter, it follows from Theorem 2.4.

For the inductive step, fix $k \geq 3$. Due to the asymptotic nature of (5.4.74), it is enough to construct the functions in $\{f_{k,n}\}_{n=1}^{\infty}$ for n larger than a certain constant of our choosing. As a starting point, the inductive hypothesis gives an explicit family $\{f_{k-2,n}\}_{n=1}^{\infty}$ in which $f_{k-2,n}\colon \{0,1\}^n \to \{0,1\}$ has threshold degree

$$\deg_{\pm}(f_{k-2,n}) = \Omega\left(n^{\frac{k-3}{k-1}} \cdot (\log n)^{-\frac{1}{k-1}\lceil\frac{k-4}{2}\rceil\lfloor\frac{k-4}{2}\rfloor}\right) \tag{5.4.75}$$

and is computable by a monotone Boolean circuit of size $n^{O(1)}$ and depth $k-2$. We view the circuit for $f_{k-2,n}$ as composed of $k-2$ layers of alternating gates, where without loss of generality the bottom layer consists of AND gates. This last property can be forced by using $\neg f_{k-2,n}(\neg x_1, \neg x_2, \ldots, \neg x_n)$ instead of $f_{k-2,n}(x_1, x_2, \ldots, x_n)$, which interchanges the circuit's AND and OR gates without affecting the threshold degree, circuit depth, or circuit size.

Now, let $c > 0$ be the absolute constant from Theorem 5.29. For every N larger that a certain constant, we apply Theorem 5.30 with

$$\begin{align}
n &= \left\lceil N^{\frac{k-1}{k+1}} (\log N)^{-\frac{1}{k+1}\lceil\frac{k-4}{2}\rceil\lfloor\frac{k-4}{2}\rfloor-\frac{2(k-1)}{k+1}} \cdot \frac{c}{100} \right\rceil, \tag{5.4.76}\\
m &= \left\lceil N^{\frac{2}{k+1}} (\log N)^{\frac{1}{k+1}\lceil\frac{k-4}{2}\rceil\lfloor\frac{k-4}{2}\rfloor-\frac{4}{k+1}} \right\rceil, \tag{5.4.77}\\
f &= f_{k-2,n}, \tag{5.4.78}\\
I_0 &= (0,\infty), \tag{5.4.79}\\
I_1 &= (-\infty,0), \tag{5.4.80}\\
I_* &= (-\infty,\infty) \tag{5.4.81}
\end{align}$$

to obtain a function $H_N\colon \{0,1\}^N \to \{0,1\}^n$ such that the composition $F_N = f_{k-2,n} \circ H_N$ has threshold degree

$$\begin{align}
\deg_\pm(F_N) &\geq \min\left\{cm\deg_\pm(f_{k-2,n}), \frac{cN}{50m\log^2(n+m)} - n\right\}\log(n+m)\\
&= \Theta\left(N^{\frac{k-1}{k+1}} (\log N)^{-\frac{1}{k+1}\lceil\frac{k-4}{2}\rceil\lfloor\frac{k-4}{2}\rfloor-\frac{k-3}{k+1}}\right)\\
&= \Theta\left(N^{\frac{k-1}{k+1}} (\log N)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right), \tag{5.4.82}
\end{align}$$

where the second step uses (5.4.75)–(5.4.77). Moreover, Theorem 5.30 ensures that H_N is computable by an AND-OR-AND circuit of polynomial size and bottom fan-in $O(\log N)$. The bottom layer of $f_{k-2,n}$ consists of AND gates, which can be merged with the top layer of H_N to produce a circuit for $F_N = f_{k-2,n} \circ H_N$ of depth $(k-2) + 3 - 1 = k$.

We have thus constructed, for some constant N_0, a family of functions $\{F_N\}_{N=N_0}^{\infty}$ in which each $F_N\colon \{0,1\}^N \to \{0,1\}$ has threshold degree (5.4.82) and is computable by a Boolean circuit of polynomial size, depth k, and bottom fan-in $O(\log N)$. Now, take the circuit for F_N and replace the negated inputs in it with N new, unnegated inputs. The resulting monotone circuit on $2N$ variables computes F_N as a subfunction and therefore has threshold degree at least that of F_N. This completes the inductive step. $\square$

Using the pattern matrix method, we now lift the previous theorem to multiparty communication complexity.

THEOREM 5.35. *Let $k \geq 3$ be a fixed integer. Let $\ell \colon \mathbb{N} \to \mathbb{N}$ be a given function. Then there is an (explicitly given) family $\{F_n\}_{n=1}^{\infty}$, where $F_n \colon (\{0,1\}^n)^{\ell(n)} \to \{0,1\}$ is an $\ell(n)$-party communication problem with discrepancy*

$$\operatorname{disc}(F_n) \leq 2\exp\left(-\Omega\left(\left(\frac{n}{4^{\ell(n)}\ell(n)^2}\right)^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right)\right) \tag{5.4.83}$$

and communication complexity

$$\mathsf{PP}(F_n) = \Omega\left(\left(\frac{n}{4^{\ell(n)}\ell(n)^2}\right)^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right). \tag{5.4.84}$$

Moreover, F_n is computable by a Boolean circuit of polynomial size and depth $k+2$ in which the bottom three layers have fan-in $O(\log n)$, $O(4^{\ell(n)}\ell(n)^2)$, and $\ell(n)$, in that order. In particular, if $\ell(n) = O(1)$, then F_n is computable by a Boolean circuit of polynomial size, depth k, and bottom fan-in $O(\log n)$.

Proof. Theorem 5.34 constructs a family of functions $\{f_n\}_{n=1}^{\infty}$, where $f_n \colon \{0,1\}^n \to \{0,1\}$ has threshold degree

$$\deg_{\pm}(f_n) = \Omega\left(n^{\frac{k-1}{k+1}} \cdot (\log n)^{-\frac{1}{k+1}\lceil\frac{k-2}{2}\rceil\lfloor\frac{k-2}{2}\rfloor}\right) \tag{5.4.85}$$

and is computable by a Boolean circuit of polynomial size, depth k, and bottom fan-in $O(\log n)$. Now, let $c > 0$ be the absolute constant from Theorem 3.10. For any given

n, define

$$F_n = \begin{cases} \mathrm{AND}_{\ell(n)} & \text{if } n \leq 2m, \\ f_{\lfloor n/m \rfloor} \circ \mathrm{NOR}_m \circ \mathrm{AND}_{\ell(n)} & \text{otherwise,} \end{cases}$$

where $m = 2\lceil c4^{\ell(n)}\ell(n)^2 \rceil$. Then the discrepancy bound (3.10) is trivial for $n \leq 2m$, and follows from (5.4.85) and Theorem 3.10 for $n > 2m$. The lower bound (5.4.84) on the communication complexity of F_n with weakly unbounded error is now immediate by the discrepancy method (Corollary 3.8).

It remains to examine the circuit complexity of F_n. Since f_n is computable by a circuit of polynomial size, depth k, and bottom fan-in $O(\log n)$, it follows that F_n is computable by a circuit of polynomial size and depth $k+2$ in which the bottom three levels have fan-in $O(\log n)$, $O(4^{\ell(n)}\ell(n)^2)$, and $\ell(n)$, in that order. This means that for $\ell(n) = O(1)$, any gate of the bottom four levels can be computed by a circuit of polynomial size, depth 2, and bottom fan-in $O(\log n)$, which in turn yields a circuit for F_n of polynomial size, depth $(k+2) - 4 + 2 = k$, and bottom fan-in $O(\log n)$. □

Theorems 5.34 and 5.35 settle Theorems 5.1 and 5.4, respectively, from the introduction.

5.5. The sign-rank of $\mathbf{AC}^0$

We now turn to the second main result of this chapter, a near-linear lower bound on the sign-rank of constant-depth circuits. To start with, we show that our smoothing technique from Theorem 5.27 already gives an exponential lower bound on the sign-rank of $\mathbf{AC}^0$. Specifically, we prove in Section 5.5.1 that the Minsky–Papert function $\mathrm{MP}_{n^{1/3}}$ has $\exp(-O(n^{1/3}))$-smooth threshold degree $\Omega(n^{1/3})$, which by Theorem 3.11

immediately implies an $\exp(\Omega(n^{1/3}))$ lower bound on the sign-rank of an $\mathbf{AC}^0$ circuit family of depth 3. This result was originally obtained, with a longer and more demanding proof, by Razborov and Sherstov [**106**].

To obtain a near-optimal lower bound of $\exp(\Omega(n^{1-\varepsilon}))$, we use a completely different approach. It is based on the notion of *local smoothness* and is unrelated to the threshold degree analysis. In Section 5.5.2, we define local smoothness and record basic properties of locally smooth functions. In Sections 5.5.3 and 5.5.4, we develop techniques for manipulating locally smooth functions to achieve desired global behavior, without the manipulations being detectable by low-degree polynomials. To apply this machinery to constant-depth circuits, we design in Section 5.5.5 a locally smooth dual polynomial for the Minsky–Papert function. We use this dual object in Section 5.5.6 to prove an amplification theorem for *smooth* threshold degree. We apply the amplification theorem iteratively in Section 5.5.7 to construct, for any $\varepsilon > 0$, a constant-depth circuit with $\exp(-n^{1-\varepsilon})$-smooth threshold degree $\Omega(n^{1-\varepsilon})$. Finally, we present our main result on the sign-rank of $\mathbf{AC}^0$ in Section 5.5.8.

In the remainder of this section, we adopt the following additional notation. For an arbitrary subset X of Euclidean space, we write $\operatorname{diam} X = \sup_{x,x' \in X} |x - x'|$, with the convention that $\operatorname{diam} \varnothing = 0$. For a vector $x \in \mathbb{Z}^n$ and a natural number d, we let $B_d(x) = \{v \in \mathbb{Z}^n : |x - v| \leq d\}$ denote the set of *integer-valued* vectors within distance d of x. For all x,

$$|B_d(x)| = |B_d(0)| \leq 2^d \binom{n+d}{d}, \tag{5.5.1}$$

where the binomial coefficient corresponds to the number of *nonnegative* integer vectors of weight at most d. Finally, for vectors $u, v \in \mathbb{N}^n$, we define $\operatorname{cube}(u, v)$ to be the smallest Cartesian product of integer intervals that contains both u and v. Specifically,

$$\begin{aligned}\operatorname{cube}(u,v) &= \{w \in \mathbb{N}^n : \min\{u_i,v_i\} \leq w_i \leq \max\{u_i,v_i\} \text{ for all } i\}\\ &= \prod_{i=1}^{n} \{\min\{u_i,v_i\}, \min\{u_i,v_i\}+1, \ldots, \max\{u_i,v_i\}\}.\end{aligned}$$

5.5.1. A simple lower bound for depth 3.

We start by presenting a new proof of Razborov and Sherstov's exponential lower bound [**106**] on the sign-rank of $\mathbf{AC}^0$. More precisely, we prove the following stronger result that was not known before.

THEOREM 5.36. *There is a constant* $0 < c < 1$ *such that for all positive integers* m *and* r,

$$\deg_{\pm}(\mathrm{MP}_{m,r}, 12^{-m-1}) \geq \min\{m, c\sqrt{r}\}.$$

Theorem 5.36 is asymptotically optimal, and it is the first lower bound on the smooth threshold degree of the Minsky–Papert function. As we will discuss shortly, this theorem implies an $\exp(\Omega(n^{1/3}))$ lower bound on the sign-rank of $\mathbf{AC}^0$. In addition, we will use Theorem 5.36 as the base case in the inductive proof of Theorem 5.3.

Proof of Theorem 5.36. It is well-known [**93, 99, 137**] that for some constant $c > 0$ and all r, any real polynomial $p\colon \{0,1\}^r \to \mathbb{R}$ with $\|p - \mathrm{OR}_r\|_\infty \leq 0.49$ has degree at least $c\sqrt{r}$. By linear programming duality [**122**, Theorem 2.5], this approximation-theoretic fact is equivalent to the existence of a function $\psi\colon \{0,1\}^m \to \mathbb{R}$ with

$$\psi(0) > 0.49, \tag{5.5.2}$$

$$\|\psi\|_1 = 1, \tag{5.5.3}$$

$$\operatorname{orth}\psi \geq c\sqrt{r}. \tag{5.5.4}$$

The rest of the proof is a reprise of Section 5.4.2. To begin with, property (5.5.3) makes it possible to view $|\psi|$ as a probability distribution on $\{0,1\}^r$. Let μ_0, μ_1, μ_2 be the probability distributions induced by $|\psi|$ on the sets $\{0^r\}, \{x \neq 0^r : \psi(x) < 0\}$, and $\{x \neq 0^r : \psi(x) > 0\}$, respectively. It is clear from (5.5.2) that the negative part of ψ is a multiple of μ_1, whereas the positive part of ψ is a nonnegative linear combination of μ_0 and μ_2. Moreover, it follows from $\langle \psi, 1 \rangle = 0$ and $\|\psi\|_1 = 1$ that the positive and negative parts of ψ both have ℓ_1-norm 1/2. Summarizing,

$$\psi = \frac{1-\delta}{2}\mu_0 - \frac{1}{2}\mu_1 + \frac{\delta}{2}\mu_2 \tag{5.5.5}$$

for some $0 \leq \delta \leq 1$. In view of (5.5.2), we infer the more precise bound

$$0 \leq \delta < \frac{1}{50}. \tag{5.5.6}$$

Let υ be the uniform probability distribution on $\{0,1\}^r \setminus \{0^r\}$. We define

$$\lambda_0 = \mu_0, \tag{5.5.7}$$

$$\lambda_1 = \frac{2}{3(1-\delta)}\mu_1 + \left(1 - \frac{2}{3(1-\delta)}\right)\upsilon, \tag{5.5.8}$$

$$\lambda_2 = \frac{2\delta}{1-\delta}\mu_2 + \left(1 - \frac{2\delta}{1-\delta}\right)\upsilon. \tag{5.5.9}$$

It is clear from (5.5.6) that λ_1 and λ_2 are convex combinations of υ, μ_1, μ_2 and therefore are probability distributions with support

$$\operatorname{supp} \lambda_i \subseteq \{0,1\}^r \setminus \{0^r\}, \qquad i = 1, 2, \tag{5.5.10}$$

whereas

$$\operatorname{supp} \lambda_0 = \{0^r\} \tag{5.5.11}$$

by definition. Moreover, (5.5.6) implies that

$$\lambda_i \geq \frac{1}{4}\upsilon, \qquad\qquad i = 1, 2. \tag{5.5.12}$$

The defining equations (5.5.7)–(5.5.9) further imply that

$$\frac{2}{3}\lambda_0 + \frac{1}{3}\lambda_2 - \lambda_1 = \frac{4}{3(1-\delta)}\psi,$$

which along with (5.5.4) gives

$$\operatorname{orth}\left(\frac{2}{3}\lambda_0 + \frac{1}{3}\lambda_2 - \lambda_1\right) \geq c\sqrt{r}. \tag{5.5.13}$$

With this work behind us, define

$$\Lambda = \frac{1}{2}\left(\frac{2}{3}\lambda_0 + \frac{1}{3}\lambda_2\right)^{\otimes m} - \frac{1}{2}\left(-\frac{1}{3}\lambda_0 + \frac{1}{3}\lambda_2\right)^{\otimes m} + \frac{1}{2}\lambda_1^{\otimes m}.$$

Multiplying out the tensor products in the definition of Λ and collecting like terms, we obtain

$$\begin{aligned}
\Lambda &= \frac{1}{2} \sum_{\substack{S \subseteq \{1,2,\ldots,m\} \\ S \neq \varnothing}} \frac{2^{|S|} - (-1)^{|S|}}{3^m} \lambda_0^{\otimes S} \cdot \lambda_2^{\otimes \overline{S}} + \frac{1}{2} \lambda_1^{\otimes m} && (5.5.14) \\
&\geq \frac{1}{4} \sum_{\substack{S \subseteq \{1,2,\ldots,m\} \\ S \neq \varnothing}} \frac{2^{|S|}}{3^m} \lambda_0^{\otimes S} \cdot \lambda_2^{\otimes \overline{S}} + \frac{1}{2} \lambda_1^{\otimes m} \\
&\geq \frac{1}{4} \sum_{\substack{S \subseteq \{1,2,\ldots,m\} \\ S \neq \varnothing}} \frac{2^{|S|}}{3^m} \lambda_0^{\otimes S} \cdot \left(\frac{1}{4}\upsilon\right)^{\otimes \overline{S}} + \frac{1}{2} \left(\frac{1}{4}\upsilon\right)^{\otimes m} \\
&\geq \frac{1}{4} \sum_{S \subseteq \{1,2,\ldots,m\}} \frac{2^{|S|}}{3^m} \lambda_0^{\otimes S} \cdot \left(\frac{1}{4}\upsilon\right)^{\otimes \overline{S}} \\
&= \frac{1}{4} \left(\frac{2}{3}\lambda_0 + \frac{1}{3} \cdot \frac{1}{4}\upsilon\right)^{\otimes m} \\
&\geq \frac{1}{4} \left(\frac{1}{12 \cdot 2^r}\right)^m \mathbf{1}_{(\{0,1\}^r)^m}, && (5.5.15)
\end{aligned}$$

where the third step uses (5.5.12). In particular, Λ is a nonnegative function. We further calculate

$$\begin{aligned}
\langle \Lambda, 1 \rangle &= \frac{1}{2} \left\langle \frac{2}{3}\lambda_0 + \frac{1}{3}\lambda_2, 1 \right\rangle^m - \frac{1}{2} \left\langle -\frac{1}{3}\lambda_0 + \frac{1}{3}\lambda_2, 1 \right\rangle^m + \frac{1}{2} \langle \lambda_1, 1 \rangle^m \\
&= \frac{1}{2} \left\langle \frac{2}{3}\lambda_0 + \frac{1}{3}\lambda_2, 1 \right\rangle^m + \frac{1}{2} \langle \lambda_1, 1 \rangle^m \\
&= \frac{1}{2} + \frac{1}{2} \\
&= 1, && (5.5.16)
\end{aligned}$$

which makes Λ a probability distribution on $(\{0,1\}^r)^m$.

It remains to examine the orthogonal content of $\Lambda \cdot (-1)^{\mathrm{MP}_{m,r}}$. We have

$$\begin{aligned}
\Lambda \cdot (-1)^{\mathrm{MP}_{m,r}} &= \frac{1}{2} \sum_{\substack{S \subseteq \{1,2,\ldots,m\} \\ S \neq \varnothing}} \frac{2^{|S|} - (-1)^{|S|}}{3^m} \lambda_0^{\otimes S} \cdot \lambda_2^{\otimes \overline{S}} \cdot (-1)^{\mathrm{MP}_{m,r}} \\
&\qquad + \frac{1}{2} \lambda_1^{\otimes m} \cdot (-1)^{\mathrm{MP}_{m,r}} \\
&= \frac{1}{2} \sum_{\substack{S \subseteq \{1,2,\ldots,m\} \\ S \neq \varnothing}} \frac{2^{|S|} - (-1)^{|S|}}{3^m} \lambda_0^{\otimes S} \cdot \lambda_2^{\otimes \overline{S}} - \frac{1}{2} \lambda_1^{\otimes m} \\
&= \frac{1}{2} \left(\frac{2}{3} \lambda_0 + \frac{1}{3} \lambda_2 \right)^{\otimes m} - \frac{1}{2} \left(-\frac{1}{3} \lambda_0 + \frac{1}{3} \lambda_2 \right)^{\otimes m} - \frac{1}{2} \lambda_1^{\otimes m},
\end{aligned}$$

where the first step uses (5.5.14); the second step uses (5.5.10) and (5.5.11); and the final equality can be verified by multiplying out the tensor powers and collecting like terms. Now

$$\begin{aligned}
&\operatorname{orth}(\Lambda \cdot (-1)^{\mathrm{MP}_{m,r}}) \\
&\qquad = \min \left\{ \operatorname{orth} \left(\frac{1}{2} \left(\frac{2}{3} \lambda_0 + \frac{1}{3} \lambda_2 \right)^{\otimes m} - \frac{1}{2} \lambda_1^{\otimes m} \right), \right. \\
&\qquad\qquad\qquad \left. \operatorname{orth} \left(-\frac{1}{2} \left(-\frac{1}{3} \lambda_0 + \frac{1}{3} \lambda_2 \right)^{\otimes m} \right) \right\} \\
&\qquad \geq \min \left\{ \operatorname{orth} \left(\frac{2}{3} \lambda_0 + \frac{1}{3} \lambda_2 - \lambda_1 \right), m \operatorname{orth} \left(-\frac{1}{3} \lambda_0 + \frac{1}{3} \lambda_2 \right) \right\} \\
&\qquad \geq \min \left\{ c\sqrt{r}, m \operatorname{orth} \left(-\frac{1}{3} \lambda_0 + \frac{1}{3} \lambda_2 \right) \right\} \\
&\qquad \geq \min\{c\sqrt{r}, m\},
\end{aligned}$$

where the first step applies Proposition 5.7(i); the second step applies Proposition 5.7(ii), (iii); the third step substitutes the lower bound from (5.5.13); and the last step uses $\langle -\lambda_0 + \lambda_2, 1 \rangle = -\langle \lambda_0, 1 \rangle + \langle \lambda_2, 1 \rangle = -1 + 1 = 0$. Combining this conclusion with (5.5.15) and (5.5.16) completes the proof. □

We now lift the approximation-theoretic result just obtained to a sign-rank lower bound, reproving a result of Razborov and Sherstov [**106**].

Theorem 5.37 (Razborov and Sherstov). *Define* $F_n\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ *by*

$$F_n = \mathrm{AND}_{n^{1/3}} \circ \mathrm{OR}_{n^{2/3}} \circ \mathrm{AND}_2.$$

Then

$$\mathrm{rk}_\pm(F_n) \geq 2^{\Omega(n^{1/3})}.$$

Proof. Theorem 5.36 states that

$$\deg_\pm(\mathrm{AND}_{n^{1/3}} \circ \mathrm{OR}_{n^{2/3}}, \exp(-c'n^{1/3})) \geq c''n^{1/3}$$

for some absolute constants $c', c'' > 0$ and all n. This lower bound along with Theorem 3.11 implies that the composition

$$H_n = \mathrm{AND}_{n^{1/3}} \circ \mathrm{OR}_{n^{2/3}} \circ \mathrm{OR}_{2\lceil \exp(\frac{4c'}{c''})\rceil} \circ \mathrm{AND}_2$$

has sign-rank $\mathrm{rk}_\pm(H_n) = \exp(\Omega(n^{1/3}))$. This completes the proof because for some integer constant $c \geq 1$, each H_n is a subfunction of F_{cn}. □

5.5.2. Local smoothness. The remainder of this chapter focuses on our $\exp(\Omega(n^{1-\varepsilon}))$ lower bound on the sign-rank of $\mathbf{AC}^0$, whose proof is unrelated to the work in Section 5.5.1. Central to our approach is an analytic notion that we call *local smoothness.* Formally, let $\Phi\colon \mathbb{N}^n \to \mathbb{R}$ be a function of interest. For a subset $X \subseteq \mathbb{N}^n$ and a real number $K \geq 1$, we say that Φ is K*-smooth on* X if

$$|\Phi(x)| \leq K^{|x-x'|}|\Phi(x')| \quad \text{for all } x, x' \in X.$$

Put another way, for any two points of X at distance d, the corresponding values of Φ differ in magnitude by a factor of at most K^d. For any set X, we let $\operatorname{Smooth}(K, X)$ denote the family of functions that are smooth on X. The following proposition collects basic properties of local smoothness, to which we refer as the restriction property, scaling property, tensor property, and conical property.

Proposition 5.38. *Let $K \geq 1$ be given.*

(i) *If $\Phi \in \operatorname{Smooth}(K, X)$ and $X' \subseteq X$, then $\Phi \in \operatorname{Smooth}(K, X')$.*

(ii) *If $\Phi \in \operatorname{Smooth}(K, X)$ and $a \in \mathbb{R}$, then $a\Phi \in \operatorname{Smooth}(K, X)$.*

(iii) $\operatorname{Smooth}(K, X) \otimes \operatorname{Smooth}(K, Y) \subseteq \operatorname{Smooth}(K, X \times Y)$.

(iv) *If $\Phi, \Psi \in \operatorname{Smooth}(K, X)$ and Φ, Ψ are nonnegative on X, then* $\operatorname{cone}\{\Phi, \Psi\} \subseteq \operatorname{Smooth}(K, X)$.

Proof. Properties (i) and (ii) are immediate from the definition of K-smoothness. For (iii), fix $(x, y), (x', y') \in X \times Y$ arbitrarily. Then

$$\begin{aligned} |\Phi(x)\Psi(y)| &\leq K^{|x-x'|}|\Phi(x')|\, K^{|y-y'|}|\Psi(y')| \\ &= K^{|(x,y)-(x',y')|}|\Phi(x')\Psi(y')|, \end{aligned}$$

where the first step uses the K-smoothness of Φ and Ψ. Finally, for (iv), let a and b be nonnegative reals. Then

$$\begin{aligned} |a\Phi(x) + b\Psi(x)| &= a|\Phi(x)| + b|\Psi(x)| \\ &\leq aK^{|x-x'|}|\Phi(x')| + bK^{|x-x'|}|\Psi(x')| \\ &= K^{|x-x'|}|a\Phi(x') + b\Psi(x')| \end{aligned}$$

for all $x, x' \in X$, where the second step uses the K-smoothness of Φ and Ψ. □

We will take a special interest in locally smooth functions that are probability distributions. For our purposes, it will be sufficient to consider locally smooth distributions whose support is the Cartesian product of integer intervals. By way of notation, for an integer $n \geq 1$ and a real number $K \geq 1$, we let $\mathfrak{S}(n, K)$ denote the set of probability distributions Λ such that:

(i) Λ is supported on $\prod_{i=1}^{n}\{0, 1, 2, \ldots, r_i\}$, for some $r_1, r_2, \ldots, r_n \in \mathbb{N}$;

(ii) Λ is K-smooth on its support.

Analogous to the development in Section 5.4.1, it will be helpful to have notation for translates of distributions in $\mathfrak{S}(n, K)$. For $\Delta \geq 0$, we let $\mathfrak{S}(n, K, \Delta)$ denote the set of probability distributions $\Lambda \in \mathfrak{D}(\mathbb{N}^n)$ such that $\Lambda(t_1, \ldots, t_n) \equiv \Lambda'(t_1 - a_1, \ldots, t_n - a_n)$ for some fixed $\Lambda' \in \mathfrak{S}(n, K)$ and $a \in \mathbb{N}^n|_{\leq \Delta}$. As a special case, $\mathfrak{S}(n, K, 0) = \mathfrak{S}(n, K)$. Specializing Proposition 5.38(iii) to this context, we obtain:

PROPOSITION 5.39. *For any $n', n'', \Delta', \Delta'', K,$ one has*

$$\mathfrak{S}(n', K, \Delta') \otimes \mathfrak{S}(n'', K, \Delta'') \subseteq \mathfrak{S}(n' + n', K, \Delta' + \Delta'').$$

Proof. The only nontrivial property to verify is K-smoothness, which follows from Proposition 5.38(iii). □

5.5.3. Metric properties of locally smooth distributions. If Λ is a locally smooth distribution on $X = \prod_{i=1}^{n}\{0, 1, 2, \ldots, r_i\}$, then a moment's thought reveals that $\Lambda(x) > 0$ at every point $x \in X$. In general, local smoothness provides one with considerable control of the metric behavior of X, making it possible to prove nontrivial upper and lower bounds on $\Lambda(S)$ for various sets $S \subseteq X$. We now record two such results, as regards our work on the sign-rank on $\mathbf{AC}^0$.

PROPOSITION 5.40. *Let Λ be a probability distribution on $X = \prod_{i=1}^{n}\{0, 1, 2, \ldots, r_i\}$. Let θ and d be nonnegative integers with $\theta \geq d$. If Λ is K-smooth on $X|_{\leq\theta}$, then*

$$\Lambda(X|_{\leq\theta}) \leq K^d \binom{n+d}{d} \Lambda(X|_{\leq\theta-d}).$$

Proof. Consider an arbitrary vector $x \in X|_{\leq\theta}$. By definition, the components of x are nonnegative integers that sum to at most θ. By decreasing the components of x as needed, one can obtain a vector x' with

$$x' \in X|_{\leq\theta-d},$$

$$x' \leq x,$$

$$|x' - x| \leq d.$$

In particular, the K-smoothness of Λ implies that

$$\Lambda(x) \leq K^d \Lambda(x').$$

Summing on both sides over $x \in X|_{\leq\theta}$, we obtain

$$\begin{aligned}\Lambda(X|_{\leq\theta}) &\leq K^d \Lambda(X|_{\leq\theta-d}) \max_{x' \in X|_{\leq\theta-d}} |\{x \in X|_{\leq\theta} : x \geq x' \text{ and } |x - x'| \leq d\}| \\ &\leq K^d \Lambda(X|_{\leq\theta-d}) \max_{x' \in \mathbb{N}^n} |\{x \in \mathbb{N}^n : x \geq x' \text{ and } |x - x'| \leq d\}| \\ &= K^d \Lambda(X|_{\leq\theta-d}) \binom{n+d}{d}.\end{aligned}$$

□

PROPOSITION 5.41. *Let Λ be a probability distribution on $X = \prod_{i=1}^{n}\{0, 1, 2, \ldots, r_i\}$. Let θ and d be nonnegative integers with*

$$d < \frac{1}{2} \min \left\{ \theta, \sum_{i=1}^{n} r_i \right\}. \qquad (5.5.17)$$

If Λ is K-smooth on $X|_{\leq\theta}$, then

$$\Lambda(X|_{\leq\theta}) \leq 2^{d+1}K^{2d+1}\binom{n+d}{d}\Lambda(X|_{\leq\theta} \setminus B_d(u))$$

for every $u \in X$.

Proof. Fix $u \in X$ for the rest of the proof. If $|u| > \theta + d$, then $X|_{\leq\theta} \setminus B_d(u) = X|_{\leq\theta}$ and the statement holds trivially. In what follows, we treat the complementary case $|u| \leq \theta + d$. Here, the key is to find a vector u' with

$$|u - u'| = d + 1, \tag{5.5.18}$$

$$u' \in X|_{\leq\theta}. \tag{5.5.19}$$

The algorithm for finding u' depends on $|u|$, as follows.

(i) If $|u| > d$, decrease one or more of the components of u as needed to obtain a vector u' whose components are nonnegative integers that sum to exactly $|u| - d - 1$. Then (5.5.18) is immediate, whereas (5.5.19) follows in view of $|u| \leq \theta + d$.

(ii) If $|u| \leq d$, the analysis is more subtle. Recall that $u \in \prod_{i=1}^{n}\{0, 1, 2, \ldots, r_i\}$ and therefore $|(r_1, \ldots, r_n) - u| = \sum r_i - |u| \geq \sum r_i - d > d$, where the last step uses (5.5.17). As a result, by increasing the components of u as necessary, one can obtain a vector $u' \in \prod_{i=1}^{n}\{0, 1, 2, \ldots, r_i\}$ with $|u'| = |u| + d + 1$. Then property (5.5.18) is immediate. Property (5.5.19) follows from $|u'| = |u| + d + 1 \leq 2d + 1 < \theta + 1$, where the last step uses (5.5.17).

Now that u' has been constructed, apply the K-smoothness of Λ to conclude that for every $x \in X|_{\leq\theta} \cap B_d(u)$,

$$\begin{aligned}\Lambda(x) &\leq K^{|x-u'|}\Lambda(u') \\ &\leq K^{|x-u|+|u-u'|}\Lambda(u') \\ &\leq K^{2d+1}\Lambda(u'),\end{aligned} \tag{5.5.20}$$

where the last step uses (5.5.18). As a result,

$$\begin{aligned}\Lambda(X|_{\leq\theta}\cap B_d(u)) &\leq \left| X|_{\leq\theta}\cap B_d(u)\right| K^{2d+1}\Lambda(u') \\ &\leq |B_d(u)|\, K^{2d+1}\Lambda(u') \\ &\leq |B_d(u)|\, K^{2d+1}\Lambda(X|_{\leq\theta}\setminus B_d(u)) \\ &\leq 2^d\binom{n+d}{d} K^{2d+1}\Lambda(X|_{\leq\theta}\setminus B_d(u)), \end{aligned} \qquad (5.5.21)$$

where the first inequality is the result of summing (5.5.20) over $x\in X|_{\leq\theta}\cap B_d(u)$; the third step uses (5.5.18) and (5.5.19); and the last step applies (5.5.1). To complete the proof, add $\Lambda(X|_{\leq\theta}\setminus B_d(u))$ to both sides of (5.5.21). □

5.5.4. Weight transfer in locally smooth distributions. Locally smooth functions exhibit great plasticity. In what follows, we will show that a locally smooth function on $\prod_{i=1}^n\{0,1,2,\ldots,r_i\}$ can be modified to achieve a broad range of global metric behaviors—without the modification being detectable by low-degree polynomials. Among other things, we will be able to take any locally smooth distribution and make it globally min-smooth. Our starting point is a generalization of Lemma 5.16, which corresponds to taking $v=0^n$ in the new result.

LEMMA 5.42. *Fix points $u,v\in\mathbb{N}^n$ and a natural number $d<|u-v|$. Then there is a function $\zeta_{u,v}\colon \operatorname{cube}(u,v)\to\mathbb{R}$ such that*

$$\operatorname{supp}\zeta_{u,v}\subseteq\{u\}\cup\{x\in\operatorname{cube}(u,v): |x-v|\leq d\}, \qquad (5.5.22)$$

$$\zeta_{u,v}(u)=1, \qquad (5.5.23)$$

$$\|\zeta_{u,v}\|_1\leq 1+2^d\binom{|u-v|}{d}, \qquad (5.5.24)$$

$$\operatorname{orth}\zeta_{u,v}>d. \qquad (5.5.25)$$

Proof. Abbreviate $u^* = (|u_1 - v_1|, |u_2 - v_2|, \ldots, |u_n - v_n|)$. Lemma 5.16 constructs a function $\zeta_{u^*} \colon \mathbb{N}^n \to \mathbb{R}$ such that

$$\operatorname{supp} \zeta_{u^*} \subseteq \{u^*\} \cup \{x \in \mathbb{N}^n : x \leq u^* \text{ and } |x| \leq d\}, \tag{5.5.26}$$

$$\zeta_{u^*}(u^*) = 1, \tag{5.5.27}$$

$$\|\zeta_{u^*}\|_1 \leq 1 + 2^d \binom{|u^*|}{d}, \tag{5.5.28}$$

$$\operatorname{orth} \zeta_{u^*} > d. \tag{5.5.29}$$

Define $\zeta_{u,v} \colon \operatorname{cube}(u,v) \to \mathbb{R}$ by

$$\zeta_{u,v}(x) = \zeta_{u^*}(|x_1 - v_1|, |x_2 - v_2|, \ldots, |x_n - v_n|).$$

Then (5.5.22) and (5.5.23) are immediate from (5.5.26) and (5.5.27), respectively. Property (5.5.24) can be verified as follows:

$$\begin{aligned} \|\zeta_{u,v}\|_1 &= \sum_{x \in \operatorname{cube}(u,v)} \zeta_{u^*}(|x_1 - v_1|, |x_2 - v_2|, \ldots, |x_n - v_n|) \\ &= \sum_{\substack{w \in \mathbb{N}^n: \\ w \leq u^*}} \zeta_{u^*}(w) \\ &\leq 1 + 2^d \binom{|u^*|}{d}, \end{aligned}$$

where the last step uses (5.5.28). For (5.5.25), fix an arbitrary polynomial p of degree at most d. Then at every point $x \in \operatorname{cube}(u,v)$, we have

$$\begin{aligned} p(x) &= p((x_1 - v_1) + v_1, \ldots, (x_n - v_n) + v_n) \\ &= p(\operatorname{sgn}(u_1 - v_1)|x_1 - v_1| + v_1, \ldots, \operatorname{sgn}(u_n - v_n)|x_n - v_n| + v_n) \\ &= q(|x_1 - v_1|, \ldots, |x_n - v_n|), \end{aligned} \tag{5.5.30}$$

where q is some polynomial of degree at most d. As a result,

$$\begin{aligned}
\langle \zeta_{u,v}, p\rangle &= \sum_{x\in\mathrm{cube}(u,v)} \zeta_{u^*}(|x_1-v_1|,\ldots,|x_n-v_n|)\,p(x)\\
&= \sum_{x\in\mathrm{cube}(u,v)} \zeta_{u^*}(|x_1-v_1|,\ldots,|x_n-v_n|)\,q(|x_1-v_1|,\ldots,|x_n-v_n|)\\
&= \sum_{\substack{w\in\mathbb{N}^n:\\ w\leq u^*}} \zeta_{u^*}(w)\,q(w)\\
&= \langle \zeta_{u^*}, q\rangle\\
&= 0,
\end{aligned}$$

where the second, fourth, and fifth steps are valid by (5.5.30), (5.5.26), and (5.5.29), respectively. □

Our next result is a smooth analogue of Lemma 5.42. The smoothness offers a great deal of flexibility when using the lemma to transfer weight from one region of $\mathbb{N}^n$ to another, in a way that cannot be detected by a low-degree polynomial.

LEMMA 5.43. *Let $X = \prod_{i=1}^n \{0,1,2,\ldots,r_i\}$, where each $r_i \geq 0$ is an integer. Let θ and d be nonnegative integers with*

$$d < \frac{1}{3}\min\left\{\theta, \sum_{i=1}^n r_i\right\}.$$

Let Λ be a probability distribution on $X|_{\leq\theta}$. Suppose further that Λ is K-smooth on $X|_{\leq\theta}$. Then for every $u \in X$, there is a function $Z_u \colon \mathbb{N}^n \to \mathbb{R}$ with

$$Z_u(u) = 1, \tag{5.5.31}$$

$$\operatorname{orth} Z_u > d, \tag{5.5.32}$$

$$\|Z_u\|_1 \leq 2^d \binom{\operatorname{diam}(\{u\} \cup \operatorname{supp} \Lambda)}{d} + 1, \tag{5.5.33}$$

$$|Z_u(x)| \leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\{u\} \cup \operatorname{supp} \Lambda)}{d} \Lambda(x), \qquad x \neq u. \tag{5.5.34}$$

Proof. We have

$$\begin{aligned} 1 &= \Lambda(X|_{\leq\theta}) \\ &\leq K^d \binom{n+d}{d} \Lambda(X|_{\leq\theta-d}) \\ &\leq 2^{d+1} K^{3d+1} \binom{n+d}{d}^2 \Lambda(X|_{\leq\theta-d} \setminus B_d(u)), \end{aligned} \tag{5.5.35}$$

where the last two step apply Propositions 5.40 and 5.41, respectively.

We now move on to the construction of Z_u. For any $v \in X|_{\leq\theta-d} \setminus B_d(u)$, Lemma 5.42 gives a function $\zeta_{u,v} \colon \mathbb{N}^n \to \mathbb{R}$ with

$$\operatorname{supp} \zeta_{u,v} \subseteq X|_{\leq\theta} \cup \{u\}, \tag{5.5.36}$$

$$\zeta_{u,v}(u) = 1, \tag{5.5.37}$$

$$\operatorname{orth} \zeta_{u,v} > d, \tag{5.5.38}$$

$$\|\zeta_{u,v}\|_1 \leq 2^d \binom{|u-v|}{d} + 1. \tag{5.5.39}$$

The last inequality can be simplified as follows:

$$\begin{aligned}\|\zeta_{u,v}\|_1 &\le 2^d \binom{\operatorname{diam}(X|_{\le\theta} \cup \{u\})}{d} + 1 \\ &\le 2^d \binom{\operatorname{diam}(\{u\} \cup \operatorname{supp}\Lambda)}{d} + 1, \qquad (5.5.40)\end{aligned}$$

where the first step uses $v \in X|_{\le\theta}$, and the second step is legitimate because Λ is a K-smooth probability distribution on $X|_{\le\theta}$ and therefore $\Lambda \ne 0$ at every point of $X|_{\le\theta}$. Combining (5.5.37) and (5.5.40),

$$\|\zeta_{u,v}\|_\infty \le 2^d \binom{\operatorname{diam}(\{u\} \cup \operatorname{supp}\Lambda)}{d}. \qquad (5.5.41)$$

We define $Z_u \colon \mathbb{N}^n \to \mathbb{R}$ by

$$Z_u(x) = \frac{1}{\Lambda(X|_{\le\theta-d} \setminus B_d(u))} \sum_{v \in X|_{\le\theta-d} \setminus B_d(u)} \Lambda(v)\, \zeta_{u,v}(x),$$

which is legitimate since $\Lambda(X|_{\le\theta-d} \setminus B_d(u)) > 0$ by (5.5.35). Then properties (5.5.31), (5.5.32), and (5.5.33) for Z_u are immediate from the corresponding properties (5.5.37), (5.5.38), and (5.5.40) of $\zeta_{u,v}$.

It remains to verify (5.5.34). Fix $x \ne u$. If $x \notin X|_{\le\theta}$, then (5.5.36) implies that $Z_u(x) = 0$ and therefore (5.5.34) holds in that case. In the complementary case when

$x \in X|_{\leq\theta}$, we have

$$\begin{aligned}|Z_u(x)| &= \sum_{v \in X|_{\leq\theta-d}\setminus B_d(u)} \frac{\Lambda(v)}{\Lambda(X|_{\leq\theta-d}\setminus B_d(u))}\cdot|\zeta_{u,v}(x)| \\ &= \sum_{\substack{v \in X|_{\leq\theta-d}\setminus B_d(u):\\ |v-x|\leq d}} \frac{\Lambda(v)}{\Lambda(X|_{\leq\theta-d}\setminus B_d(u))}\cdot|\zeta_{u,v}(x)| \\ &\leq \sum_{\substack{v \in X|_{\leq\theta-d}\setminus B_d(u):\\ |v-x|\leq d}} \frac{K^d\Lambda(x)}{\Lambda(X|_{\leq\theta-d}\setminus B_d(u))}\cdot 2^d\binom{\operatorname{diam}(\{u\}\cup\operatorname{supp}\Lambda)}{d} \\ &\leq 2^d\binom{n+d}{d}\cdot\frac{K^d\Lambda(x)}{\Lambda(X|_{\leq\theta-d}\setminus B_d(u))}\cdot 2^d\binom{\operatorname{diam}(\{u\}\cup\operatorname{supp}\Lambda)}{d},\end{aligned}$$

where the first step applies the triangle inequality to the definition of Z_u; the second step uses (5.5.36) and $x \neq u$; the third step applies the K-smoothness of Λ and substitutes the bound from (5.5.41); and the final step uses (5.5.1). In view of (5.5.35), this completes the proof of (5.5.34). □

We now show how to efficiently zero out a locally smooth function on points of large Hamming weight. The modified function is pointwise close to the original and cannot be distinguished from it by any low-degree polynomial.

LEMMA 5.44. *Define $X = \prod_{i=1}^{n}\{0,1,2,\ldots,r_i\}$, where each $r_i \geq 0$ is an integer. Let θ and d be nonnegative integers with*

$$d < \frac{\theta}{3}. \tag{5.5.42}$$

Let $\Phi\colon X \to \mathbb{R}$ *be a function that is* K*-smooth on* $X|_{\leq\theta}$*, with* $\Phi|_{\leq\theta} \not\equiv 0$*. Then there is* $\tilde{\Phi}\colon X \to \mathbb{R}$ *such that*

$$\operatorname{orth}(\Phi - \tilde{\Phi}) > d, \tag{5.5.43}$$

$$\operatorname{supp} \tilde{\Phi} \subseteq X|_{\leq\theta}, \tag{5.5.44}$$

$$|\Phi - \tilde{\Phi}| \leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp}\Phi)}{d} \frac{\|\Phi|_{>\theta}\|_1}{\|\Phi|_{\leq\theta}\|_1} \cdot |\Phi| \qquad \textit{on } X|_{\leq\theta}. \tag{5.5.45}$$

Proof. If $\theta > \sum_{i=1}^n r_i$, the lemma holds trivially for $\tilde{\Phi} = \Phi$. In what follows, we treat the complementary case $\theta \leq \sum_{i=1}^n r_i$. By (5.5.42),

$$d < \frac{1}{3} \min\left\{\theta, \sum_{i=1}^n r_i\right\}.$$

Since Φ is K-smooth on $X|_{\leq\theta}$, the probability distribution Λ on $X|_{\leq\theta}$ given by $\Lambda(x) = |\Phi(x)|/\|\Phi|_{\leq\theta}\|_1$ is also K-smooth. As a result, Lemma 5.43 gives for every $u \in X$ a function $Z_u\colon X \to \mathbb{R}$ with

$$Z_u(u) = 1, \tag{5.5.46}$$

$$|Z_u(x)| \leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\{u\} \cup \operatorname{supp}\Lambda)}{d} \frac{|\Phi(x)|}{\|\Phi|_{\leq\theta}\|_1} \qquad \text{for } x \neq u, \tag{5.5.47}$$

$$\operatorname{orth} Z_u > d, \tag{5.5.48}$$

$$\operatorname{supp} Z_u \subseteq X|_{\leq\theta} \cup \{u\}. \tag{5.5.49}$$

Now define

$$\tilde{\Phi} = \Phi - \sum_{u \in X|_{>\theta}} \Phi(u) Z_u.$$

Then (5.5.43) is immediate from (5.5.48). To verify (5.5.44), fix any point $x \in X|_{>\theta}$. Then

$$\begin{aligned}\tilde{\Phi}(x) &= \Phi(x) - \sum_{u \in X|_{>\theta}} \Phi(u) Z_u(x) \\ &= \Phi(x) - \Phi(x) Z_x(x) \\ &= 0,\end{aligned}$$

where the last two steps use (5.5.49) and (5.5.46), respectively.

It remains to verify (5.5.45) on $X|_{\leq\theta}$:

$$\begin{aligned}|\Phi - \tilde{\Phi}| &\leq \sum_{\substack{u \in X|_{>\theta}: \\ \Phi(u) \neq 0}} \Phi(u)\,|Z_u| \\ &\leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp}\Phi)}{d} \sum_{\substack{u \in X|_{>\theta}: \\ \Phi(u) \neq 0}} |\Phi(u)| \cdot \frac{|\Phi|}{\|\Phi|_{\leq\theta}\|_1} \\ &= 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp}\Phi)}{d} \frac{\|\Phi|_{>\theta}\|_1}{\|\Phi|_{\leq\theta}\|_1} \cdot |\Phi|,\end{aligned}$$

where the second step uses (5.5.47). □

For technical reasons, we need a generalization of the previous lemma to functions on $\prod_{i=1}^n \{\Delta_i, \Delta_i + 1, \ldots, \Delta_i + r_i\}$ for nonnegative integers Δ_i and r_i, and further to convex combinations of such functions. We obtain these generalizations in the two corollaries that follow.

COROLLARY 5.45. *Define $X = \prod_{i=1}^n \{\Delta_i, \Delta_i + 1, \ldots, \Delta_i + r_i\}$, where all Δ_i and r_i are nonnegative integers. Let θ and d be nonnegative integers with*

$$d < \frac{1}{3}\left(\theta - \sum_{i=1}^n \Delta_i\right).$$

Let $\Phi\colon X \to \mathbb{R}$ be a function that is K-smooth on $X|_{\leq\theta}$, with $\Phi|_{\leq\theta} \not\equiv 0$. Then there is a function $\tilde{\Phi}\colon X \to \mathbb{R}$ such that

$$\operatorname{orth}(\Phi - \tilde{\Phi}) > d, \tag{5.5.50}$$

$$\operatorname{supp} \tilde{\Phi} \subseteq X|_{\leq\theta}, \tag{5.5.51}$$

$$|\Phi - \tilde{\Phi}| \leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp}\Phi)}{d} \frac{\|\Phi|_{>\theta}\|_1}{\|\Phi|_{\leq\theta}\|_1} \cdot |\Phi| \qquad \textit{on } X|_{\leq\theta.} \tag{5.5.52}$$

Proof. Abbreviate $X' = \prod_{i=1}^n \{0, 1, 2, \ldots, r_i\}$ and $\theta' = \theta - \sum_{i=1}^n \Delta_i$. In this notation,

$$d < \frac{\theta'}{3}. \tag{5.5.53}$$

Consider the function $\Phi'\colon X' \to \mathbb{R}$ given by $\Phi'(x) = \Phi(x + (\Delta_1, \Delta_2 \ldots, \Delta_n))$. Then any two points $u, v \in X'|_{\leq\theta'}$ obey

$$\begin{aligned} |\Phi'(u)| &= |\Phi(u + (\Delta_1, \Delta_2, \ldots, \Delta_n))| \\ &\leq K^{|u-v|} |\Phi(v + (\Delta_1, \Delta_2, \ldots, \Delta_n))| \\ &= K^{|u-v|} |\Phi'(v)|, \end{aligned}$$

where the second step uses the K-smoothness of Φ on $X|_{\leq\theta}$ As a result, Φ' is K-smooth on $X'|_{\leq\theta'}$. Moreover, $\|\Phi'|_{\leq\theta'}\|_1 = \|\Phi|_{\leq\theta}\|_1 > 0$. In view of (5.5.53), Lemma 5.44 gives a function a function $\tilde{\Phi}'\colon X' \to \mathbb{R}$ such that

$$\operatorname{orth}(\Phi' - \tilde{\Phi}') > d,$$

$$\operatorname{supp} \tilde{\Phi}' \subseteq X'|_{\leq\theta'},$$

and

$$|\Phi' - \tilde{\Phi}'| \leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp} \Phi')}{d} \frac{\|\Phi'|_{>\theta'}\|_1}{\|\Phi'|_{\leq\theta'}\|_1} \cdot |\Phi'|$$
$$= 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp} \Phi)}{d} \frac{\|\Phi|_{>\theta}\|_1}{\|\Phi|_{\leq\theta}\|_1} \cdot |\Phi'|$$

on $X'|_{\leq\theta'}$. As a result, (5.5.50)–(5.5.52) hold for the real-valued function $\tilde{\Phi}\colon X \to \mathbb{R}$ given by $\tilde{\Phi}(x) = \tilde{\Phi}'(x - (\Delta_1, \Delta_2, \ldots, \Delta_n))$. □

COROLLARY 5.46. *Fix integers $\Delta, d, \theta \geq 0$ and $n \geq 1$, and a real number δ, where*

$$\delta \in [0, 1),$$
$$d < \frac{1}{3}(\theta - \Delta).$$

Then for every

$$\Lambda \in \operatorname{conv}(\mathfrak{S}(n, K, \Delta) \cap \{\Lambda' \in \mathfrak{D}(\mathbb{N}^n) : \Lambda'(\mathbb{N}^n|_{>\theta}) \leq \delta\}),$$

there is a function $\tilde{\Lambda}\colon \mathbb{N}^n \to \mathbb{R}$ such that

$$\operatorname{orth}(\Lambda - \tilde{\Lambda}) > d,$$
$$\operatorname{supp} \tilde{\Lambda} \subseteq \mathbb{N}^n|_{\leq\theta} \cap \operatorname{supp} \Lambda,$$
$$|\Lambda - \tilde{\Lambda}| \leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp} \Lambda)}{d} \frac{\delta}{1-\delta} \cdot \Lambda \qquad \text{on } \mathbb{N}^n|_{\leq\theta}.$$

Proof. Write Λ out explicitly as

$$\Lambda = \sum_{i=1}^{N} \lambda_i \Lambda_i$$

for some positive reals $\lambda_1, \ldots, \lambda_N$ with $\sum \lambda_i = 1$, where $\Lambda_i \in \mathfrak{S}(n, K, \Delta)$ and $\Lambda_i(\mathbb{N}^n|_{>\theta}) \leq \delta$. Then clearly

$$\operatorname{supp} \Lambda = \bigcup_{i=1}^{n} \operatorname{supp} \Lambda_i. \tag{5.5.54}$$

For $i = 1, 2, \ldots, N$, Corollary 5.45 constructs $\tilde{\Lambda}_i \colon \mathbb{N}^n \to \mathbb{R}$ with

$$\operatorname{orth}(\Lambda_i - \tilde{\Lambda}_i) > d, \tag{5.5.55}$$

$$\operatorname{supp} \tilde{\Lambda}_i \subseteq \mathbb{N}^n|_{\leq\theta}, \tag{5.5.56}$$

$$|\Lambda_i - \tilde{\Lambda}_i| \leq 2^{3d+1} K^{4d+1} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp} \Lambda_i)}{d} \frac{\delta}{1-\delta} \cdot \Lambda_i \qquad \text{on } \mathbb{N}^n|_{\leq\theta}, \tag{5.5.57}$$

$$\operatorname{supp} \tilde{\Lambda}_i \subseteq \operatorname{supp} \Lambda_i. \tag{5.5.58}$$

In view of (5.5.54)–(5.5.58), the proof is complete by taking $\tilde{\Lambda} = \sum_{i=1}^{N} \lambda_i \tilde{\Lambda}_i$. □

Our next result uses local smoothness to achieve something completely different. Here, we show how to start with a locally smooth function and make it globally min-smooth. The new function has the same sign pointwise as the original, and cannot be distinguished from it by any low-degree polynomial. Crucially for us, the global min-smoothness can be achieved relative to any distribution on the domain.

LEMMA 5.47. *Define $X = \prod_{i=1}^{n} \{0, 1, 2, \ldots, r_i\}$, where each $r_i \geq 0$ is an integer. Let θ and d be nonnegative integers with*

$$d < \frac{1}{3} \min \left\{ \theta, \sum_{i=1}^{n} r_i \right\}.$$

Let $\Phi\colon X|_{\leq\theta} \to \mathbb{R}$ *be a function that is* K*-smooth on* $X|_{\leq\theta}$*. Then for every probability distribution* Λ^* *on* $X|_{\leq\theta}$*, there is* $\Phi^*\colon X|_{\leq\theta} \to \mathbb{R}$ *such that*

$$\operatorname{orth}(\Phi - \Phi^*) > d, \tag{5.5.59}$$

$$\|\Phi^*\|_1 \leq 2\|\Phi\|_1, \tag{5.5.60}$$

$$\Phi \cdot \Phi^* \geq 0, \tag{5.5.61}$$

$$|\Phi^*| \geq \left(2^{3d+1}K^{4d+1}\binom{n+d}{d}^3\binom{\operatorname{diam}(\operatorname{supp}\Phi)}{d}\right)^{-1}\|\Phi\|_1\,\Lambda^*. \tag{5.5.62}$$

Proof. If $\Phi \equiv 0$, the lemma holds trivially with $\Phi^* = \Phi$. In the complementary case, abbreviate

$$N = 2^{3d+1}K^{4d+1}\binom{n+d}{d}^3\binom{\operatorname{diam}(\operatorname{supp}\Phi)}{d}.$$

We will view $|\Phi|/\|\Phi\|_1$ as a probability distribution on $X|_{\leq\theta}$. By hypothesis, this probability distribution is K-smooth on $X|_{\leq\theta}$. In particular, $X|_{\leq\theta} \subseteq \operatorname{supp}|\Phi| = \operatorname{supp}\Phi$. Therefore, Lemma 5.43 gives for every $u \in X|_{\leq\theta}$ a function $Z_u\colon X|_{\leq\theta} \to \mathbb{R}$ with

$$Z_u(u) = 1, \tag{5.5.63}$$

$$\|Z_u\|_1 \leq \frac{N}{2} + 1, \tag{5.5.64}$$

$$|Z_u(x)| \leq N \cdot \frac{|\Phi(x)|}{\|\Phi\|_1}, \qquad x \neq u, \tag{5.5.65}$$

$$\operatorname{orth} Z_u > d. \tag{5.5.66}$$

Now, define $\Phi^*\colon X|_{\leq\theta} \to \mathbb{R}$ by

$$\Phi^* = \Phi + \frac{\|\Phi\|_1}{N}\sum_{u\in X|_{\leq\theta}} \widetilde{\operatorname{sign}}(\Phi(u))\Lambda^*(u)Z_u.$$

Then (5.5.59) follows directly from (5.5.66). For (5.5.60), we have:

$$\begin{aligned}
\|\Phi^*\|_1 &\leq \|\Phi\|_1 + \frac{\|\Phi\|_1}{N} \sum_{u \in X|_{\leq\theta}} \Lambda^*(u)\, \|Z_u\|_1 \\
&\leq \|\Phi\|_1 + \frac{\|\Phi\|_1}{N} \cdot \left(\frac{N}{2} + 1\right) \sum_{u \in X|_{\leq\theta}} \Lambda^*(u) \\
&= \frac{3N+2}{2N} \|\Phi\|_1 \\
&\leq 2\, \|\Phi\|_1, \qquad (5.5.67)
\end{aligned}$$

where the second step uses (5.5.64). The remaining properties (5.5.61) and (5.5.62) can be established simultaneously as follows: for every $x \in X|_{\leq\theta}$,

$$\begin{aligned}
&\widetilde{\operatorname{sign}}(\Phi(x)) \cdot \Phi^*(x) \\
&\qquad = |\Phi(x)| + \frac{\|\Phi\|_1}{N} \sum_{u \in X|_{\leq\theta}} \Lambda^*(u) Z_u(x) \\
&\qquad \geq |\Phi(x)| + \frac{\|\Phi\|_1}{N} \Lambda^*(x) Z_x(x) - \frac{\|\Phi\|_1}{N} \sum_{\substack{u \in X|_{\leq\theta}: \\ u \neq x}} \Lambda^*(u)\, |Z_u(x)| \\
&\qquad = |\Phi(x)| + \frac{\|\Phi\|_1}{N} \Lambda^*(x) - \frac{\|\Phi\|_1}{N} \sum_{\substack{u \in X|_{\leq\theta}: \\ u \neq x}} \Lambda^*(u)\, |Z_u(x)| \\
&\qquad \geq |\Phi(x)| + \frac{\|\Phi\|_1}{N} \Lambda^*(x) - \frac{\|\Phi\|_1}{N} \cdot N \cdot \frac{|\Phi(x)|}{\|\Phi\|_1} \sum_{\substack{u \in X|_{\leq\theta}: \\ u \neq x}} \Lambda^*(u) \\
&\qquad = |\Phi(x)| + \frac{\|\Phi\|_1}{N} \Lambda^*(x) - |\Phi(x)|\, (1 - \Lambda^*(x)) \\
&\qquad \geq \frac{\|\Phi\|_1}{N} \Lambda^*(x), \qquad (5.5.68)
\end{aligned}$$

where the third and fourth steps use (5.5.63) and (5.5.65), respectively. □

5.5.5. A locally smooth dual polynomial for MP. As Sections 5.5.2–5.5.4 show, local smoothness implies several useful metric and analytic properties. To tap into this resource, we now construct a locally smooth dual polynomial for the Minsky–Papert function. It is helpful to view this new result as a counterpart of Theorem 5.27 from our analysis of the threshold degree of $\mathbf{AC}^0$. The new proof is considerably more technical because local smoothness is a delicate property to achieve.

THEOREM 5.48. *For some absolute constant* $0 < c < 1$ *and all positive integers* m, r, R *with* $r \leq R$, *there are probability distributions* Λ_0 *and* Λ_1 *such that*

$$\operatorname{supp}\Lambda_0 = (\mathrm{MP}^*_{m,R})^{-1}(0), \tag{5.5.69}$$

$$\operatorname{supp}\Lambda_1 = (\mathrm{MP}^*_{m,R})^{-1}(1), \tag{5.5.70}$$

$$\operatorname{orth}(\Lambda_0 - \Lambda_1) \geq \min\{m, c\sqrt{r}\}, \tag{5.5.71}$$

$$\frac{\Lambda_0 + \Lambda_1}{2} \in \operatorname{Smooth}\left(\frac{m}{c}, \{0, 1, 2, \ldots, R\}^m\right), \tag{5.5.72}$$

$$\Lambda_0, \Lambda_1 \in \operatorname{conv}\left(\left\{\lambda \in \mathfrak{S}\left(1, \frac{1}{c}, 1\right) : \right.\right.$$
$$\left.\left.\lambda(t) \leq \frac{1}{c(t+1)^2\, 2^{ct/\sqrt{r}}} \textit{ for } t \in \mathbb{N}\right\}^{\otimes m}\right). \tag{5.5.73}$$

Our proof of Theorem 5.48 repeatedly employs the following simple but useful criterion for K-smoothness: a probability distribution λ is K-smooth on an integer interval $I = \{i, i+1, i+2, \ldots, j\}$ if and only if the probabilities of any two *consecutive* integers in I are within a factor of K.

Proof of Theorem 5.48. Abbreviate $\varepsilon = 1/6$. For some absolute constants $c', c'' \in (0,1)$, Lemma 5.26 constructs probability distributions $\lambda_0, \lambda_1, \lambda_2$ such that

$$\operatorname{supp}\lambda_0 = \{0\}, \tag{5.5.74}$$

$$\operatorname{supp}\lambda_i = \{1, 2, \ldots, R\}, \qquad i = 1, 2, \tag{5.5.75}$$

$$\lambda_i(t) \in \left[\frac{c'}{t^2\, 2^{c''t/\sqrt{r}}},\ \frac{1}{c't^2\, 2^{c''t/\sqrt{r}}}\right], \qquad i = 1, 2; \quad t = 1, 2, \ldots, R, \tag{5.5.76}$$

$$\operatorname{orth}((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2 - \lambda_1) \geq c'\sqrt{r}. \tag{5.5.77}$$

We infer that

$$\lambda_0 \in \mathfrak{S}(1, K), \tag{5.5.78}$$

$$\lambda_1 \in \mathfrak{S}(1, K, 1), \tag{5.5.79}$$

$$\lambda_2 \in \mathfrak{S}(1, K, 1), \tag{5.5.80}$$

$$(1-\varepsilon)\lambda_0 + \varepsilon\lambda_2 \in \mathfrak{S}(1, K), \tag{5.5.81}$$

$$\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}\lambda_1 \in \mathfrak{S}(1, Km) \tag{5.5.82}$$

for some large constant $K = K(c', c'') \geq 1$. Indeed, (5.5.78) is trivial since λ_0 is the single-point distribution on the origin; (5.5.79) holds because by (5.5.75) and (5.5.76), the probabilities of any pair of consecutive integers in $\operatorname{supp}\lambda_1 = \{1, 2, \ldots, R\}$ are the same up to a constant factor; and (5.5.80)–(5.5.82) can be seen analogously, by comparing the probabilities of any pair of consecutive integers. Combining (5.5.78)–(5.5.82) with Proposition 5.39, we obtain

$$\{\lambda_0, \lambda_1, \lambda_2\}^{\otimes m} \subseteq \mathfrak{S}(m, K, m), \tag{5.5.83}$$

$$((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)^{\otimes m} \in \mathfrak{S}(m, K), \tag{5.5.84}$$

$$\left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}\lambda_1\right)^{\otimes m} \in \mathfrak{S}(m, Km). \tag{5.5.85}$$

The proof centers around the dual objects $\Psi_1, \Psi_2\colon \{0,1,2,\ldots,R\}^m \to \mathbb{R}$ given by

$$\Psi_1 = \left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}\lambda_1\right)^{\otimes m} - 2\lambda_1^{\otimes m}$$

and

$$\begin{aligned}\Psi_2 = 2((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)^{\otimes m} - 2(-\varepsilon\lambda_0 + \varepsilon\lambda_2)^{\otimes m}\\ - \left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)\right)^{\otimes m}.\end{aligned}$$

The next four claims establish key properties of Ψ_1 and Ψ_2. □

Claim 5.49. Ψ_1 *satisfies*

$$\operatorname{pos}\Psi_1 \in \operatorname{cone}(\{\lambda_0,\lambda_1\}^{\otimes m} \setminus \{\lambda_1^{\otimes m}\}), \tag{5.5.86}$$

$$\operatorname{neg}\Psi_1 \in \operatorname{cone}\{\lambda_1^{\otimes m}\}, \tag{5.5.87}$$

$$\frac{1}{5}|\Psi_1| \leq \left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}\lambda_1\right)^{\otimes m} \leq |\Psi_1|. \tag{5.5.88}$$

Claim 5.50. Ψ_2 *satisfies*

$$\operatorname{pos}\Psi_2 \in \operatorname{cone}(\{\lambda_0,\lambda_2\}^{\otimes m} \setminus \{\lambda_2^{\otimes m}\}), \tag{5.5.89}$$

$$\operatorname{neg}\Psi_2 \in \operatorname{cone}\{\lambda_2^{\otimes m}\}, \tag{5.5.90}$$

$$\frac{1}{3}|\Psi_2| \leq ((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)^{\otimes m} \leq 3|\Psi_2|. \tag{5.5.91}$$

Claim 5.51. Ψ_1 *and* Ψ_2 *satisfy*

$$\operatorname{supp}(\operatorname{pos}\Psi_i) = (\mathrm{MP}^*_{m,R})^{-1}(0), \qquad i = 1,2, \tag{5.5.92}$$

$$\operatorname{supp}(\operatorname{neg}\Psi_i) = (\mathrm{MP}^*_{m,R})^{-1}(1), \qquad i = 1,2. \tag{5.5.93}$$

Claim 5.52. $\operatorname{orth}(\Psi_1 + \Psi_2) \geq \min\{m, c'\sqrt{r}\}$.

Proof. We will settle Claims 5.49–5.52 shortly, once we complete the main proof. Define

$$\Lambda_0 = \frac{2}{\|\Psi_1\|_1 + \|\Psi_2\|_1} \operatorname{pos}(\Psi_1 + \Psi_2),$$
$$\Lambda_1 = \frac{2}{\|\Psi_1\|_1 + \|\Psi_2\|_1} \operatorname{neg}(\Psi_1 + \Psi_2),$$

where the denominators are nonzero by (5.5.88). We proceed to verify the properties required of Λ_0 and Λ_1 in the theorem statement.

Support. Recall from Claim 5.51 that the positive parts of Ψ_1 and Ψ_2 are supported on $(\mathrm{MP}^*_{m,R})^{-1}(0)$. Therefore, the positive part of $\Psi_1 + \Psi_2$ is supported on $(\mathrm{MP}^*_{m,R})^{-1}(0)$ as well, which in turn implies that

$$\operatorname{supp} \Lambda_0 = (\mathrm{MP}^*_{m,R})^{-1}(0). \tag{5.5.94}$$

Analogously, Claim 5.51 states that the negative parts of Ψ_1 and Ψ_2 are supported on $(\mathrm{MP}^*_{m,R})^{-1}(1)$. As a result, the negative part of $\Psi_1 + \Psi_2$ is also supported on $(\mathrm{MP}^*_{m,R})^{-1}(1)$, whence

$$\operatorname{supp} \Lambda_1 = (\mathrm{MP}^*_{m,R})^{-1}(1). \tag{5.5.95}$$

Orthogonality. The defining equations for Λ_0 and Λ_1 imply that

$$\Lambda_0 - \Lambda_1 = \frac{2}{\|\Psi_1\|_1 + \|\Psi_2\|_1} (\Psi_1 + \Psi_2),$$

which along with Claim 5.52 forces

$$\operatorname{orth}(\Lambda_0 - \Lambda_1) \geq \min\{m, c'\sqrt{r}\}. \tag{5.5.96}$$

Nonnegativity and norm. By definition, Λ_0 and Λ_1 are nonnegative functions. We calculate

$$\begin{aligned}\|\Lambda_0\|_1 - \|\Lambda_1\|_1 &= \langle \Lambda_0, 1\rangle - \langle \Lambda_1, 1\rangle \\ &= \langle \Lambda_0 - \Lambda_1, 1\rangle \\ &= 0, \end{aligned} \tag{5.5.97}$$

where the first step uses the nonnegativity of Λ_0 and Λ_1, and the last step applies (5.5.96). In addition,

$$\begin{aligned}\|\Lambda_0\|_1 + \|\Lambda_1\|_1 &= \frac{2}{\|\Psi_1\|_1 + \|\Psi_2\|_1}(\|\operatorname{pos}(\Psi_1 + \Psi_2)\|_1 + \|\operatorname{neg}(\Psi_1 + \Psi_2)\|_1) \\ &= \frac{2}{\|\Psi_1\|_1 + \|\Psi_2\|_1}\|\Psi_1 + \Psi_2\|_1 \\ &= 2, \end{aligned} \tag{5.5.98}$$

where the last step uses Claim 5.51. A consequence of (5.5.97) and (5.5.98) is that $\|\Lambda_0\|_1 = \|\Lambda_1\|_1 = 1$, which makes Λ_0 and Λ_1 probability distributions. In view of (5.5.94) and (5.5.95), we conclude that

$$\Lambda_i \in \mathfrak{D}((\mathrm{MP}^*_{m,R})^{-1}(i)), \qquad\qquad i = 0, 1. \tag{5.5.99}$$

In particular,

$$\frac{\Lambda_0 + \Lambda_1}{2} \in \mathfrak{D}(\{0, 1, 2, \dots, R\}^m). \tag{5.5.100}$$

Smoothness. We have

$$\begin{aligned}\frac{\Lambda_0+\Lambda_1}{2} &= \frac{|\Psi_1+\Psi_2|}{\|\Psi_1\|_1+\|\Psi_2\|_1}\\ &= \frac{1}{\|\Psi_1\|_1+\|\Psi_2\|_1}\,|\Psi_1| + \frac{1}{\|\Psi_1\|_1+\|\Psi_2\|_1}\,|\Psi_2|, \qquad (5.5.101)\end{aligned}$$

where the first step follows from the defining equations for Λ_0 and Λ_1, and the second step uses Claim 5.51. Inequality (5.5.88) shows that at every point, $|\Psi_1|$ is within a factor of 5 of the tensor product $(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}\lambda_1)^{\otimes m}$, which by (5.5.85) is Km-smooth on its support. It follows that $|\Psi_1|$ is $5Km$-smooth on $\{0,1,2,\ldots,R\}^m$. By an analogous argument, (5.5.91) and (5.5.84) imply that $|\Psi_2|$ is $3K$-smooth (and hence also $5Km$-smooth) on $\{0,1,2,\ldots,R\}^m$. Now (5.5.101) shows that $\frac{1}{2}(\Lambda_0+\Lambda_1)$ is a conical combination of two nonnegative $5Km$-smooth functions on $\{0,1,2,\ldots,R\}^m$. By Proposition 5.38(iv),

$$\frac{\Lambda_0+\Lambda_1}{2} \in \mathrm{Smooth}(5Km, \{0,1,2,\ldots,R\}^m). \qquad (5.5.102)$$

Having examined the convex combination $\frac{\Lambda_0+\Lambda_1}{2}$, we now turn to the individual distributions Λ_0 and Λ_1. We have

$$\begin{aligned}\Lambda_0 &= \frac{2}{\|\Psi_1\|_1+\|\Psi_2\|_1}\,\mathrm{pos}(\Psi_1+\Psi_2)\\ &= \frac{2}{\|\Psi_1\|_1+\|\Psi_2\|_1}(\mathrm{pos}(\Psi_1)+\mathrm{pos}(\Psi_2))\\ &\in \mathrm{cone}(\{\lambda_0,\lambda_1,\lambda_2\}^{\otimes m}),\end{aligned}$$

where the first equation restates the definition of Λ_0, the second step applies (5.5.92), and the last step uses (5.5.86) and (5.5.89). Analogously,

$$\begin{aligned}\Lambda_1 &= \frac{2}{\|\Psi_1\|_1 + \|\Psi_2\|_1}\operatorname{neg}(\Psi_1+\Psi_2)\\ &= \frac{2}{\|\Psi_1\|_1 + \|\Psi_2\|_1}(\operatorname{neg}(\Psi_1)+\operatorname{neg}(\Psi_2))\\ &\in \operatorname{cone}(\{\lambda_1^{\otimes m}, \lambda_2^{\otimes m}\}),\end{aligned}$$

where the first equation restates the definition of Λ_1, the second step applies (5.5.93), and the last step uses (5.5.87) and (5.5.90). Thus, Λ_0 and Λ_1 are conical combinations of probability distributions in $\{\lambda_0, \lambda_1, \lambda_2\}^{\otimes m}$. Since Λ_0 and Λ_1 are themselves probability distributions, we conclude that

$$\Lambda_0, \Lambda_1 \in \operatorname{conv}(\{\lambda_0, \lambda_1, \lambda_2\}^{\otimes m}).$$

By (5.5.74)–(5.5.76),

$$\lambda_i(t) \leq \frac{1}{c'''(t+1)^2\, 2^{c'''t/\sqrt{r}}} \qquad (t \in \mathbb{N};\ i = 0, 1, 2)$$

for some constant $c''' > 0$. The last two equations along with (5.5.78)–(5.5.80) yield

$$\Lambda_0, \Lambda_1 \in \operatorname{conv}\left(\left\{\lambda \in \mathfrak{S}(1, K, 1) : \lambda(t) \leq \frac{1}{c'''(t+1)^2\, 2^{c'''t/\sqrt{r}}} \text{ for } t \in \mathbb{N}\right\}^{\otimes m}\right). \qquad (5.5.103)$$

Now (5.5.94)–(5.5.96), (5.5.102), and (5.5.103) imply (5.5.69)–(5.5.73) for a small enough constant $c > 0$. □

We now settle the four claims made in the proof of Theorem 5.48.

Proof of Claim 5.49. Multiplying out the tensor product in the definition of Ψ_1 and collecting like terms, we obtain

$$\begin{aligned}\Psi_1 = &-\left(2-\left(\frac{m}{m+1}\right)^m\right)\lambda_1^{\otimes m}\\ &+\sum_{\substack{S\subseteq\{1,2,\ldots,m\}\\ S\neq\varnothing}}\left(\frac{1}{m+1}\right)^{|S|}\left(\frac{m}{m+1}\right)^{m-|S|}\lambda_0^{\otimes S}\cdot\lambda_1^{\otimes\overline{S}}.\end{aligned}\tag{5.5.104}$$

Recall from (5.5.74) and (5.5.75) that λ_0 and λ_1 are supported on $\{0\}$ and $\{1,2,\ldots,R\}$, respectively. Therefore, the right-hand side of (5.5.104) is the sum of 2^m nonzero functions whose supports are pairwise disjoint. Now (5.5.86) and (5.5.87) follow directly from (5.5.104). One further obtains that

$$\begin{aligned}|\Psi_1| = &\left(2-\left(\frac{m}{m+1}\right)^m\right)\lambda_1^{\otimes m}\\ &+\sum_{\substack{S\subseteq\{1,2,\ldots,m\}\\ S\neq\varnothing}}\left(\frac{1}{m+1}\right)^{|S|}\left(\frac{m}{m+1}\right)^{m-|S|}\lambda_0^{\otimes S}\cdot\lambda_1^{\otimes\overline{S}}.\end{aligned}$$

From first principles,

$$\begin{aligned}\left(\frac{1}{m+1}\lambda_0+\frac{m}{m+1}\lambda_1\right)^{\otimes m} = &\left(\frac{m}{m+1}\right)^m\lambda_1^{\otimes m}\\ &+\sum_{\substack{S\subseteq\{1,2,\ldots,m\}\\ S\neq\varnothing}}\left(\frac{1}{m+1}\right)^{|S|}\left(\frac{m}{m+1}\right)^{m-|S|}\lambda_0^{\otimes S}\cdot\lambda_1^{\otimes\overline{S}}.\end{aligned}$$

Comparing the right-hand sides of the last two equations settles (5.5.88). □

Proof of Claim 5.50. Multiplying out the tensor powers in the definition of Ψ_2 and collecting like terms, we obtain

$$\Psi_2 = -\left(\frac{m}{m+1}\right)^m \varepsilon^m \lambda_2^{\otimes m} + \sum_{\substack{S\subseteq\{1,2,\ldots,m\}\\ S\neq\varnothing}} a_{|S|}\, \lambda_0^{\otimes S}\cdot\lambda_2^{\otimes \overline{S}}, \tag{5.5.105}$$

where the coefficients $a_1, a_2, \ldots, a_m$ are given by

$$\begin{aligned} a_i &= \left(2(1-\varepsilon)^i\varepsilon^{m-i} - 2(-1)^i\varepsilon^m - \left(1-\frac{\varepsilon m}{m+1}\right)^i\left(\frac{\varepsilon m}{m+1}\right)^{m-i}\right) \\ &= (1-\varepsilon)^i\varepsilon^{m-i}\left(2 - 2\left(\frac{-\varepsilon}{1-\varepsilon}\right)^i - \left(\frac{m+1-\varepsilon m}{(m+1)(1-\varepsilon)}\right)^i\left(\frac{m}{m+1}\right)^{m-i}\right) \\ &\in \left[\frac{1}{3}(1-\varepsilon)^i\varepsilon^{m-i},\, 3(1-\varepsilon)^i\varepsilon^{m-i}\right]. \end{aligned} \tag{5.5.106}$$

As in the proof of the previous claim, recall from (5.5.74) and (5.5.75) that λ_0 and λ_2 have disjoint support. Therefore, the right-hand side of (5.5.105) is the sum of 2^m nonzero functions whose supports are pairwise disjoint. Now (5.5.89) and (5.5.90) are immediate from (5.5.106). The disjointness of the supports of the summands on the right-hand side of (5.5.105) also implies that

$$|\Psi_2| = \left(\frac{m}{m+1}\right)^m \varepsilon^m \lambda_0^{\otimes m} + \sum_{\substack{S\subseteq\{1,2,\ldots,m\}\\ S\neq\varnothing}} |a_{|S|}|\, \lambda_0^{\otimes S}\cdot\lambda_2^{\otimes \overline{S}}.$$

In view of (5.5.106), we conclude that $|\Psi_2|$ coincides up to a factor of 3 with the function

$$\sum_{S\subseteq\{1,2,\ldots,m\}} (1-\varepsilon)^{|S|}\varepsilon^{m-|S|}\lambda_0^{\otimes S}\cdot\lambda_2^{\otimes \overline{S}} = ((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)^{\otimes m}.$$

This settles (5.5.91) and completes the proof. □

Proof of Claim 5.51. Recall from (5.5.74) and (5.5.75) that $\operatorname{supp}\lambda_0 = \{0\}$ and $\operatorname{supp}\lambda_1 = \operatorname{supp}\lambda_2 = \{1, 2, \ldots, R\}$. In this light, (5.5.86)–(5.5.88) imply

$$\begin{aligned}
\operatorname{supp}(\operatorname{pos}\Psi_1) &\subseteq (\mathrm{MP}^*_{m,R})^{-1}(0),\\
\operatorname{supp}(\operatorname{neg}\Psi_1) &\subseteq (\mathrm{MP}^*_{m,R})^{-1}(1),\\
\operatorname{supp}(\Psi_1) &= (\mathrm{MP}^*_{m,R})^{-1}(0) \cup (\mathrm{MP}^*_{m,R})^{-1}(1),
\end{aligned}$$

respectively. Analogously, (5.5.89)–(5.5.91) imply

$$\begin{aligned}
\operatorname{supp}(\operatorname{pos}\Psi_2) &\subseteq (\mathrm{MP}^*_{m,R})^{-1}(0),\\
\operatorname{supp}(\operatorname{neg}\Psi_2) &\subseteq (\mathrm{MP}^*_{m,R})^{-1}(1),\\
\operatorname{supp}(\Psi_2) &= (\mathrm{MP}^*_{m,R})^{-1}(0) \cup (\mathrm{MP}^*_{m,R})^{-1}(1).
\end{aligned}$$

Since the support of each Ψ_i is the disjoint union of the supports of its positive and negative parts, (5.5.92) and (5.5.93) follow. □

Proof of Claim 5.52. Write $\Psi_1 + \Psi_2 = A + B + C$, where

$$\begin{aligned}
A &= \left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}\lambda_1\right)^{\otimes m} - \left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)\right)^{\otimes m},\\
B &= 2((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)^{\otimes m} - 2\lambda_1^{\otimes m},\\
C &= -2(-\varepsilon\lambda_0 + \varepsilon\lambda_2)^{\otimes m}.
\end{aligned}$$

As a result, Proposition 5.7(i) guarantees that

$$\operatorname{orth}(\Psi_1 + \Psi_2) \geq \min\{\operatorname{orth} A, \operatorname{orth} B, \operatorname{orth} C\}. \tag{5.5.107}$$

We have

$$\begin{aligned}\operatorname{orth} A &\geq \operatorname{orth}\left(\left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}\lambda_1\right)\right.\\ &\qquad\qquad \left.-\left(\frac{1}{m+1}\lambda_0 + \frac{m}{m+1}((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2)\right)\right)\\ &= \operatorname{orth}\left(-\frac{m}{m+1}((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2 - \lambda_1)\right)\\ &\geq c'\sqrt{r}, \qquad (5.5.108)\end{aligned}$$

where the first step uses Proposition 5.7(iii), and the last step is a restatement of (5.5.77). Analogously,

$$\begin{aligned}\operatorname{orth} B &\geq \operatorname{orth}(((1-\varepsilon)\lambda_0 + \varepsilon\lambda_2) - \lambda_1)\\ &\geq c'\sqrt{r}, \qquad (5.5.109)\end{aligned}$$

where the first and second steps use Proposition 5.7(iii) and (5.5.77), respectively. Finally,

$$\begin{aligned}\operatorname{orth} C &= \operatorname{orth}((-\varepsilon\lambda_0 + \varepsilon\lambda_2)^{\otimes m})\\ &= m \operatorname{orth}(-\varepsilon\lambda_0 + \varepsilon\lambda_2)\\ &\geq m, \qquad (5.5.110)\end{aligned}$$

where the second step applies Proposition 5.7(ii), and the third step is valid because $\langle -\varepsilon\lambda_0 + \varepsilon\lambda_2, 1\rangle = -\varepsilon\langle\lambda_0, 1\rangle + \varepsilon\langle\lambda_2, 1\rangle = -\varepsilon + \varepsilon = 0$. By (5.5.107)–(5.5.110), the proof is complete. □

5.5.6. An amplification theorem for smooth threshold degree. We have reached the technical centerpiece of our sign-rank analysis, an amplification theorem for smooth threshold degree. This result is considerably stronger than the amplification theorems for threshold degree in Section 5.4.3, which does not preserve smoothness. We prove the new amplification theorem by manipulating locally smooth distributions to achieve the desired global behavior, an approach unrelated to our work in Section 5.4.3. A detailed statement of our result follows.

Theorem 5.53. *There is an absolute constant $C \geq 1$ such that*

for all:

positive integers n, m, r, R, θ with $R \geq r$ and $\theta \geq Cnm \log(2nm)$;

real numbers $\gamma \in [0,1]$;

functions $f\colon \{0,1\}^n \to \{0,1\}$;

probability distributions Λ^* on $\{0,1,2,\ldots,R\}^{mn}|_{\leq\theta}$; and

positive integers d with

$$d \leq \frac{1}{C} \min\left\{ m \deg_{\pm}(f,\gamma),\ \sqrt{r}\deg_{\pm}(f,\gamma),\ \frac{\theta}{\sqrt{r}\log(2nmR)} \right\}, \tag{5.5.111}$$

one has:

$$\operatorname{orth}((-1)^{f\circ \mathrm{MP}^*_{m,R}} \cdot \Lambda) \geq d, \tag{5.5.112}$$

$$\Lambda \geq \gamma \cdot (CnmR)^{-8d}\, \Lambda^* \tag{5.5.113}$$

for some $\Lambda \in \mathfrak{D}(\{0,1,2,\ldots,R\}^{mn}|_{\leq\theta})$.

Proof. Let $0 < c < 1$ be the constant from Theorem 5.48. Take $C \geq 1/c$ to be a sufficiently large absolute constant. By hypothesis,

$$\theta \geq Cnm \log(2nm). \tag{5.5.114}$$

Abbreviate

$$\begin{aligned} X &= \{0, 1, 2, \ldots, R\}^{nm}, \\ \delta &= 2^{-c\theta/(2\sqrt{r})}. \end{aligned} \tag{5.5.115}$$

The following inequalities are straightforward to verify:

$$d < \frac{1}{3}\min\{\theta - nm, nmR\}, \tag{5.5.116}$$

$$\theta \geq \frac{8enm(1 + \ln(nm))}{c}, \tag{5.5.117}$$

$$\frac{2^{3d+1}}{c^{4d+1}}\binom{n+d}{d}^3\binom{nmR}{d}\frac{\delta}{1-\delta} < \frac{1}{2}, \tag{5.5.118}$$

$$2^{3d+1}\left(\frac{3m}{c}\right)^{4d+1}\binom{n+d}{d}^3\binom{nmR}{d} \leq \frac{(CnmR)^{8d}}{4}. \tag{5.5.119}$$

For example, (5.5.116) holds because $d \leq nm/C$ by (5.5.111) and $\theta \geq Cnm\log(2nm)$ by (5.5.114). Inequalities (5.5.117)–(5.5.119) follow analogously from (5.5.111) and (5.5.114) for a large enough constant C. The rest of the proof splits neatly into four major steps.

STEP 1: KEY DISTRIBUTIONS. Theorem 5.48 provides probability distributions Λ_0 and Λ_1 such that

$$\begin{aligned}
&\operatorname{supp}\Lambda_i = (\mathrm{MP}^*_{m,R})^{-1}(i), && i = 0,1, && (5.5.120)\\
&\operatorname{orth}(\Lambda_0 - \Lambda_1) \geq \min\{m, c\sqrt{r}\}, && && (5.5.121)\\
&\frac{\Lambda_0+\Lambda_1}{2} \in \operatorname{Smooth}\left(\frac{m}{c}, \{0,1,2,\ldots,R\}^m\right), && && (5.5.122)\\
&\Lambda_0,\Lambda_1 \in \operatorname{conv}\left(\left\{\lambda \in \mathfrak{G}\left(1,\frac{1}{c},1\right) :\right.\right. && && \\
&\qquad\qquad \left.\left.\lambda(t) \leq \frac{1}{c(t+1)^2\, 2^{ct/\sqrt{r}}} \text{ for } t \in \mathbb{N}\right\}^{\otimes m}\right). && && (5.5.123)
\end{aligned}$$

Consider the probability distributions

$$\Lambda_z = \bigotimes_{i=1}^{n} \Lambda_{z_i}, \qquad\qquad z \in \{0,1\}^n.$$

Then

$$\begin{aligned}
\Lambda_z &\in \operatorname{conv}\left(\left\{\lambda \in \mathfrak{G}\left(1,\frac{1}{c},1\right) : \lambda(t) \leq \frac{1}{c(t+1)^2\, 2^{ct/\sqrt{r}}} \text{ for } t \in \mathbb{N}\right\}^{\otimes mn}\right)\\
&\subseteq \operatorname{conv}\left(\mathfrak{G}\left(1,\frac{1}{c},1\right)^{\otimes mn} \cap \right.\\
&\qquad\qquad \left.\left\{\lambda \in \mathfrak{D}(\mathbb{N}) : \lambda(t) \leq \frac{1}{c(t+1)^2\, 2^{ct/\sqrt{r}}} \text{ for } t \in \mathbb{N}\right\}^{\otimes mn}\right)\\
&\subseteq \operatorname{conv}\left(\mathfrak{G}\left(1,\frac{1}{c},1\right)^{\otimes mn} \cap \left\{\Lambda \in \mathfrak{D}(\mathbb{N}^{mn}) : \Lambda(\mathbb{N}^{nm}|_{>\theta}) \leq 2^{-c\theta/(2\sqrt{r})}\right\}\right)\\
&\subseteq \operatorname{conv}\left(\mathfrak{G}\left(1,\frac{1}{c},1\right)^{\otimes mn} \cap \{\Lambda \in \mathfrak{D}(\mathbb{N}^{mn}) : \Lambda(\mathbb{N}^{nm}|_{>\theta}) \leq \delta\}\right)\\
&\subseteq \operatorname{conv}\left(\mathfrak{G}\left(nm,\frac{1}{c},nm\right) \cap \{\Lambda \in \mathfrak{D}(\mathbb{N}^{mn}) : \Lambda(\mathbb{N}^{nm}|_{>\theta}) \leq \delta\}\right), \qquad (5.5.124)
\end{aligned}$$

where the first step uses (5.2.1) and (5.5.123); the third step is valid by (5.5.117) and Lemma 5.20; the fourth step is a substitution from (5.5.115); and the last step is an application of Proposition 5.39.

Step 2: Restricting the support. By (5.5.116), (5.5.124), and Corollary 5.46, there is a real function $\tilde{\Lambda}_z \colon \mathbb{N}^{nm} \to \mathbb{R}$ such that

$$\operatorname{orth}(\Lambda_z - \tilde{\Lambda}_z) > d, \tag{5.5.125}$$

$$\operatorname{supp} \tilde{\Lambda}_z \subseteq \mathbb{N}^{nm}|_{\leq\theta}, \tag{5.5.126}$$

$$\operatorname{supp} \tilde{\Lambda}_z \subseteq \operatorname{supp} \Lambda_z, \tag{5.5.127}$$

and

$$|\Lambda_z - \tilde{\Lambda}_z| \leq \frac{2^{3d+1}}{c^{4d+1}} \binom{n+d}{d}^3 \binom{\operatorname{diam}(\operatorname{supp} \Lambda_z)}{d} \frac{\delta}{1-\delta} \cdot \Lambda_z \qquad \text{on } \mathbb{N}^{nm}|_{\leq\theta}.$$

In view of (5.5.118) and $\operatorname{diam}(\operatorname{supp} \Lambda_z) \leq nmR$, the last equation simplifies to

$$|\Lambda_z - \tilde{\Lambda}_z| \leq \frac{1}{2}\Lambda_z \quad \text{on } \mathbb{N}^{nm}|_{\leq\theta}. \tag{5.5.128}$$

Properties (5.5.126) and (5.5.128) imply that $\tilde{\Lambda}_z$ is a nonnegative function, which along with (5.5.125) and Proposition 5.10 implies that $\tilde{\Lambda}_z$ is a probability distribution. Combining this fact with (5.5.120), (5.5.126), and (5.5.127) gives

$$\tilde{\Lambda}_z \in \mathfrak{D}\left(\mathbb{N}^{nm}|_{\leq\theta} \cap \prod_{i=1}^{n} (\mathrm{MP}^*_{m,R})^{-1}(z_i)\right), \qquad z \in \{0,1\}^n. \tag{5.5.129}$$

In particular, the $\tilde{\Lambda}_z$ are supported on disjoint sets of inputs.

Step 3: Ensuring min-smoothness. Recall from (5.5.129) that each of the probability distributions $\tilde{\Lambda}_z$ is supported on a subset of $X|_{\leq\theta}$. Consider the function

$\Phi\colon X|_{\leq\theta} \to \mathbb{R}$ given by

$$\Phi = 2^{-n} \sum_{z\in\{0,1\}^n} (-1)^{f(z)} \tilde{\Lambda}_z.$$

Again by (5.5.129), the support of $\tilde{\Lambda}_z$ is contained in $\prod_{i=1}^n (\mathrm{MP}^*_{m,R})^{-1}(z_i)$. This means in particular that $f \circ \mathrm{MP}^*_{m,R} = f(z)$ on the support of $\tilde{\Lambda}_z$, whence

$$(-1)^{f(z)} \tilde{\Lambda}_z = (-1)^{f\circ \mathrm{MP}^*_{m,R}} \cdot \tilde{\Lambda}_z \tag{5.5.130}$$

everywhere on $X|_{\leq\theta}$. Making this substitution in the defining equation for Φ, we find that

$$(-1)^{f\circ \mathrm{MP}^*_{m,R}} \cdot \Phi \geq 0. \tag{5.5.131}$$

The fact that the $\tilde{\Lambda}_z$ are supported on pairwise disjoint sets of inputs forces

$$|\Phi| = 2^{-n} \sum_{z\in\{0,1\}^n} \tilde{\Lambda}_z \tag{5.5.132}$$

and in particular

$$\|\Phi\|_1 = 1. \tag{5.5.133}$$

We now examine the smoothness of Φ. For this, consider the probability distribution

$$\Lambda = 2^{-n} \sum_{z\in\{0,1\}^n} \Lambda_z. \tag{5.5.134}$$

Comparing equations (5.5.132) and (5.5.134) term by term and using the upper bound (5.5.128), we find that $|\Lambda - |\Phi|| \leq \frac{1}{2}\Lambda$ on $X|_{\leq\theta}$. Equivalently,

$$\frac{1}{2}\Lambda \leq |\Phi| \leq \frac{3}{2}\Lambda \quad \text{on } X|_{\leq\theta.} \tag{5.5.135}$$

But

$$\begin{aligned}\Lambda &= \left(\frac{1}{2}\Lambda_0 + \frac{1}{2}\Lambda_1\right)^{\otimes n} \\ &\in \operatorname{Smooth}\left(\frac{m}{c}, \{0,1,2,\ldots,R\}^m\right)^{\otimes n} \\ &\subseteq \operatorname{Smooth}\left(\frac{m}{c}, \{0,1,2,\ldots,R\}^{mn}\right), \qquad (5.5.136)\end{aligned}$$

where the last two steps are valid by (5.5.122) and Proposition 5.38(iii), respectively. Combining (5.5.135) and (5.5.136), we conclude that Φ is $(3m/c)$-smooth on $X|_{\leq\theta}$. As a result, (5.5.116) and Lemma 5.47 provide a function $\Phi^*\colon X|_{\leq\theta} \to \mathbb{R}$ with

$$\operatorname{orth}(\Phi - \Phi^*) > d, \qquad (5.5.137)$$

$$\|\Phi^*\|_1 \leq 2\|\Phi\|_1, \qquad (5.5.138)$$

$$\Phi \cdot \Phi^* \geq 0, \qquad (5.5.139)$$

$$|\Phi^*| \geq \left(2^{3d+1}\left(\frac{3m}{c}\right)^{4d+1}\binom{n+d}{d}^3\binom{\operatorname{diam}(\operatorname{supp}\Phi)}{d}\right)^{-1}\|\Phi\|_1\,\Lambda^*. \qquad (5.5.140)$$

In view of (5.5.133), the second property simplifies to

$$\|\Phi^*\|_1 \leq 2. \qquad (5.5.141)$$

Recall that on $X|_{\leq\theta}$, the function Φ is $(3m/c)$-smooth and not identically zero. Therefore, Φ must be nonzero at every point of $X|_{\leq\theta}$, which includes the support of Φ^*. As a result, (5.5.131) and (5.5.139) imply that

$$(-1)^{f\circ\mathrm{MP}_{m,R}} \cdot \Phi^* \geq 0. \qquad (5.5.142)$$

Finally, using $\operatorname{diam}(\operatorname{supp}\Phi) \leq nmR$ along with the bounds (5.5.119) and (5.5.133), we can restate (5.5.140) as

$$|\Phi^*| \geq 4(CnmR)^{-8d}\Lambda^*. \qquad (5.5.143)$$

Step 4: The final construction. By the definition of smooth threshold degree, there is a probability distribution μ on $\{0,1\}^n$ such that

$$\operatorname{orth}((-1)^f \cdot \mu) \geq \deg_{\pm}(f, \gamma), \tag{5.5.144}$$

$$\mu(z) \geq \gamma \cdot 2^{-n}, \qquad z \in \{0,1\}^n. \tag{5.5.145}$$

Define

$$\Phi_{\text{final}} = \sum_{z \in \{0,1\}^n} \mu(z)(-1)^{f(z)} \tilde{\Lambda}_z - \gamma\Phi + \gamma\Phi^*.$$

The right-hand side is a linear combination of functions on $X|_{\leq\theta}$, whence

$$\operatorname{supp}(\Phi_{\text{final}}) \subseteq X|_{\leq\theta}. \tag{5.5.146}$$

Moreover,

$$\begin{aligned}\|\Phi_{\text{final}}\|_1 &\leq \sum_{z \in \{0,1\}^n} \mu(z)\|\tilde{\Lambda}_z\|_1 + \gamma\|\Phi\|_1 + \gamma\|\Phi^*\|_1 \\ &\leq 1 + 3\gamma \\ &\leq 4, \end{aligned} \tag{5.5.147}$$

where the first step applies the triangle inequality, and the second step uses (5.5.129), (5.5.133) and (5.5.141). Continuing,

$$\begin{aligned}
&(-1)^{f\circ \mathrm{MP}^*_{m,R}} \cdot \Phi_{\text{final}} \\
&\quad = (-1)^{f\circ \mathrm{MP}^*_{m,R}} \cdot \left(\sum_{z\in\{0,1\}^n} (\mu(z) - \gamma 2^{-n})(-1)^{f(z)}\tilde{\Lambda}_z + \gamma\Phi^* \right) \\
&\quad = \sum_{z\in\{0,1\}^n} (\mu(z) - \gamma 2^{-n})(-1)^{f\circ \mathrm{MP}^*_{m,R}} \cdot (-1)^{f(z)}\tilde{\Lambda}_z + \gamma(-1)^{f\circ \mathrm{MP}^*_{m,R}} \cdot \Phi^* \\
&\quad = \sum_{z\in\{0,1\}^n} (\mu(z) - \gamma 2^{-n})\tilde{\Lambda}_z + \gamma|\Phi^*| && (5.5.148) \\
&\quad \geq \gamma|\Phi^*| \\
&\quad \geq 4\gamma(CnmR)^{-8d}\Lambda^*, && (5.5.149)
\end{aligned}$$

where the first step applies the definition of Φ; the third step uses (5.5.130) and (5.5.142); the fourth step follows from (5.5.145); and the fifth step substitutes the lower bound from (5.5.143). Now

$$\Phi_{\text{final}} \not\equiv 0 \tag{5.5.150}$$

follows from (5.5.148) if $\gamma = 0$, and from (5.5.149) if $\gamma > 0$.

It remains to examine the orthogonal content of Φ_{final}. For this, write

$$\Phi_{\text{final}} = \sum_{z\in\{0,1\}^n} \mu(z)(-1)^{f(z)}\Lambda_z + \sum_{z\in\{0,1\}^n} \mu(z)(-1)^{f(z)}(\tilde{\Lambda}_z - \Lambda_z) \\ + \gamma(\Phi^* - \Phi).$$

Then

$$\begin{aligned}
\operatorname{orth}(\Phi_{\text{final}}) &\geq \min\left\{\operatorname{orth}\left(\sum_{z\in\{0,1\}^n} \mu(z)(-1)^{f(z)}\Lambda_z\right),\right. \\
&\qquad\qquad\qquad \left.\min_z\{\operatorname{orth}(\tilde{\Lambda}_z - \Lambda_z)\},\;\; \operatorname{orth}(\Phi^* - \Phi)\right\} \\
&\geq \min\left\{\operatorname{orth}\left(\sum_{z\in\{0,1\}^n} \mu(z)(-1)^{f(z)}\Lambda_z\right), d\right\} \\
&\geq \min\left\{\operatorname{orth}\left(\sum_{z\in\{0,1\}^n} \mu(z)(-1)^{f(z)}\bigotimes_{i=1}^n \Lambda_{z_i}\right), d\right\} \\
&\geq \min\left\{\operatorname{orth}(\mu\cdot(-1)^f)\operatorname{orth}(\Lambda_1 - \Lambda_0), d\right\} \\
&\geq \min\{\deg_\pm(f,\gamma)\min\{m, c\sqrt{r}\}, d\} \\
&= d, \qquad\qquad (5.5.151)
\end{aligned}$$

where the first step applies Proposition 5.7(i); the second step follows from (5.5.125) and (5.5.137); the third step is valid by the definition of Λ_z; the fourth step applies Corollary 5.9; the fifth step substitutes the lower bounds from (5.5.121) and (5.5.144); and the final step uses (5.5.111).

To complete the proof, let

$$\Lambda = \frac{\Phi_{\text{final}}}{\|\Phi_{\text{final}}\|_1}\cdot(-1)^{f\circ\text{MP}^*_{m,R}},$$

where the right-hand side is well-defined by (5.5.150). Then $\|\Lambda\|_1 = 1$ by definition. Moreover, (5.5.146) and (5.5.149) guarantee that Λ is a nonnegative function with support contained in $X|_{\leq\theta}$, so that $\Lambda \in \mathfrak{D}(X|_{\leq\theta})$. The orthogonality property (5.5.112) follows from (5.5.151), whereas the min-smoothness property (5.5.113) follows from (5.5.147) and (5.5.149). □

We now translate the new amplification theorem from $\mathbb{N}^n|_{\leq\theta}$ to the hypercube, using the input transformation scheme of Theorem 5.23.

THEOREM 5.54. *Let $C \geq 1$ be the absolute constant from Theorem* 5.53. *Fix positive integers n, m, θ with $\theta \geq Cnm\log(2nm)$. Then there is an* (*explicitly given*) *transformation $H\colon \{0,1\}^{6\theta\lceil\log(nm+1)\rceil} \to \{0,1\}^n$, computable by an AND-OR-AND circuit of polynomial size with bottom fan-in at most* $6\lceil\log(nm+1)\rceil$, *such that*

$$\deg_{\pm}(f \circ H, \gamma\theta^{-24d}) \geq d\lceil\log(nm+1)+1\rceil, \tag{5.5.152}$$

$$\deg_{\pm}(f \circ \neg H, \gamma\theta^{-24d}) \geq d\lceil\log(nm+1)+1\rceil \tag{5.5.153}$$

for all Boolean functions $f\colon \{0,1\}^n \to \{0,1\}$, all real numbers $\gamma \in [0,1]$, and all positive integers

$$d \leq \frac{1}{C}\min\left\{m\deg_{\pm}(f,\gamma), \frac{\theta}{4m\log\theta}\right\}.$$

Proof. Negating a function's input bits has no effect on its γ-smooth threshold degree for any $0 \leq \gamma \leq 1$, so that $f(x_1, x_2, \ldots, x_n)$ and $f(\neg x_1, \neg x_2, \ldots, \neg x_n)$ both have γ-smooth threshold degree $\deg_{\pm}(f,\gamma)$. Therefore, proving (5.5.152) for all f will also settle (5.5.153) for all f. In what follows, we focus on the former.

Theorem 5.23 constructs an explicit surjection $G\colon \{0,1\}^N \to \mathbb{N}^{nm}|_{\leq\theta}$ on $N = 6\theta\lceil\log(nm+1)\rceil$ variables with the following two properties:

(i) for every coordinate $i = 1, 2, \ldots, nm$, the mapping $x \mapsto \mathrm{OR}^*_\theta(G(x)_i)$ is computable by a DNF formula of size $(nm\theta)^{O(1)} = \theta^{O(1)}$ with bottom fan-in at most $6\lceil\log(nm+1)\rceil$;

(ii) for any polynomial p, the map $v \mapsto \mathbf{E}_{G^{-1}(v)}\, p$ is a polynomial on $\mathbb{N}^{nm}|_{\leq\theta}$ of degree at most $(\deg p)/\lceil\log(nm+1)+1\rceil$.

Consider the composition $F = (f \circ \mathrm{MP}^*_{m,\theta}) \circ G$. Then

$$\begin{aligned} F &= (f \circ (\mathrm{AND}_m \circ \mathrm{OR}^*_\theta)) \circ G \\ &= f \circ ((\underbrace{\mathrm{AND}_m \circ \mathrm{OR}^*_\theta, \ldots, \mathrm{AND}_m \circ \mathrm{OR}^*_\theta}_{n}) \circ G), \end{aligned}$$

which by property (i) of G means that F is the composition of f and an AND-OR-AND circuit H of size $(nm\theta)^{O(1)} = \theta^{O(1)}$ and bottom fan-in $6\lceil \log(nm+1) \rceil$. Hence, the proof will be complete once we show that

$$\deg_\pm(F, \gamma\theta^{-24d}) \geq d\lceil \log(nm+1) + 1 \rceil. \tag{5.5.154}$$

Define $r = m^2$ and $R = \max\{\theta, r\}$, and consider the probability distribution on $\{0, 1, 2, \ldots, R\}^{nm}|_{\leq\theta} = \mathbb{N}^{nm}|_{\leq\theta}$ given by $\Lambda^*(v) = |G^{-1}(v)|/2^N$. Then Theorem 5.53 constructs a probability distribution Λ on $\mathbb{N}^{nm}|_{\leq\theta}$ such that

$$\operatorname{orth}((-1)^{f \circ \mathrm{MP}^*_{m,R}} \cdot \Lambda) \geq d, \tag{5.5.155}$$

$$\Lambda \geq \gamma\theta^{-24d}\, \Lambda^*. \tag{5.5.156}$$

In view of $R \geq \theta$, inequality (5.5.155) can be restated as

$$\operatorname{orth}((-1)^{f \circ \mathrm{MP}^*_{m,\theta}} \cdot \Lambda) \geq d. \tag{5.5.157}$$

Define

$$\lambda = \sum_{v \in \mathbb{N}^{nm}|_{\leq\theta}} \Lambda(v) \cdot \frac{\mathbf{1}_{G^{-1}(v)}}{|G^{-1}(v)|},$$

where $\mathbf{1}_{G^{-1}(v)}$ denotes as usual the characteristic function of the set $G^{-1}(v)$. Clearly, λ is a probability distribution on $\{0,1\}^N$. Moreover,

$$\begin{aligned}\lambda &\geq \gamma\theta^{-24d} \sum_{v\in\mathbb{N}^{nm}|_{\leq\theta}} \Lambda^*(v)\cdot\frac{\mathbf{1}_{G^{-1}(v)}}{|G^{-1}(v)|}\\ &= \gamma\theta^{-24d} \sum_{v\in\mathbb{N}^{nm}|_{\leq\theta}} \frac{|G^{-1}(v)|}{2^N}\cdot\frac{\mathbf{1}_{G^{-1}(v)}}{|G^{-1}(v)|}\\ &= \gamma\theta^{-24d}\cdot\frac{\mathbf{1}_{\{0,1\}^N}}{2^N}, \end{aligned} \tag{5.5.158}$$

where the first two steps use (5.5.156) and the definition of Λ^*, respectively.

Finally, we examine the orthogonal content of $(-1)^F\cdot\lambda$. Let $p\colon \mathbb{R}^N\to\mathbb{R}$ be any polynomial of degree less than $d\lceil\log(nm+1)+1\rceil$. Then by property (ii) of G, the mapping $p^*\colon v\mapsto \mathbf{E}_{G^{-1}(v)}\, p$ is a polynomial on $\mathbb{N}^{nm}|_{\leq\theta}$ of degree less than d. As a result,

$$\begin{aligned}\langle(-1)^F\cdot\lambda,p\rangle &= \langle(-1)^{(f\circ\mathrm{MP}^*_{m,\theta})\circ G}\cdot\lambda,p\rangle\\ &= \sum_{v\in\mathbb{N}^{nm}|_{\leq\theta}}\ \sum_{G^{-1}(v)} (-1)^{(f\circ\mathrm{MP}^*_{m,\theta})\circ G}\cdot\lambda\cdot p\\ &= \sum_{v\in\mathbb{N}^{nm}|_{\leq\theta}} (-1)^{(f\circ\mathrm{MP}^*_{m,\theta})(v)} \sum_{G^{-1}(v)}\lambda\cdot p\\ &= \sum_{v\in\mathbb{N}^{nm}|_{\leq\theta}} (-1)^{(f\circ\mathrm{MP}^*_{m,\theta})(v)}\Lambda(v) \mathop{\mathbf{E}}_{G^{-1}(v)} p\\ &= \langle(-1)^{f\circ\mathrm{MP}^*_{m,\theta}}\cdot\Lambda,p^*\rangle\\ &= 0,\end{aligned}$$

where the last step uses (5.5.157) and $\deg p^* < d$. We conclude that $\mathrm{orth}((-1)^F\cdot\lambda)\geq d\lceil\log(nm+1)+1\rceil$, which along with (5.5.158) settles (5.5.154). □

5.5.7. The smooth threshold degree of $\mathbf{AC^0}$. We now construct, for any $\varepsilon > 0$, a constant-depth circuit $f\colon \{0,1\}^n \to \{0,1\}$ with $\exp(-n^{1-\varepsilon})$-smooth threshold degree $\Omega(n^{1-\varepsilon})$. This result may find applications in future work, in addition to its use in this chapter to obtain a lower bound on the sign-rank of $\mathbf{AC}^0$. The proof proceeds by induction, with the amplification theorem for smooth threshold degree (Theorem 5.54) applied repeatedly to construct increasingly harder circuits. To simplify the exposition, we isolate the inductive step in the following lemma.

Lemma 5.55. *Let $f\colon \{0,1\}^n \to \{0,1\}$ be a Boolean circuit of size s, depth d, and smooth threshold degree*

$$\deg_{\pm}\left(f, \exp\left(-c'\cdot\frac{n^{1-\alpha}}{\log^{\beta} n}\right)\right) \geq c''\cdot\frac{n^{1-\alpha}}{\log^{\beta} n},$$

for some real numbers $\alpha \in [0,1]$, $\beta \geq 0$, and $c', c'' > 0$. Then f can be transformed in polynomial time into a Boolean circuit $F\colon \{0,1\}^N \to \{0,1\}$ on $N = \Theta(n^{1+\alpha}\log^{2+\beta} n)$ variables that has size $s + N^{O(1)}$, depth at most $d+3$, bottom fan-in $O(\log n)$, and smooth threshold degree

$$\deg_{\pm}\left(F, \exp\left(-C'\cdot\frac{N^{\frac{1}{1+\alpha}}}{\log^{\frac{1-\alpha+\beta}{1+\alpha}} N}\right)\right) \geq C''\cdot\frac{N^{\frac{1}{1+\alpha}}}{\log^{\frac{1-\alpha+\beta}{1+\alpha}} N}, \tag{5.5.159}$$

where $C', C'' > 0$ are real numbers that depend on c', c'' only. Moreover, if the circuit for f is monotone with AND gates at the bottom, then the depth of F is at most $d+2$.

Proof. Let $C \geq 1$ be the absolute constant from Theorem 5.53. Apply Theorem 5.54 with

$$m = \lceil n^{\alpha}\log^{\beta} n\rceil,$$

$$\theta = \lceil Cmn\log(2nm)\rceil,$$

$$\gamma = \exp\left(-c'\cdot\frac{n^{1-\alpha}}{\log^{\beta} n}\right)$$

to obtain a function $H\colon \{0,1\}^N \to \{0,1\}^n$ on $N = \Theta(n^{1+\alpha}\log^{2+\beta} n)$ variables such that the composition $F = f \circ H$ satisfies (5.5.159) for some $C', C'' > 0$ that depend only on c', c'', and furthermore H is computable by an AND-OR-AND circuit of polynomial size and bottom fan-in $O(\log N)$. Clearly, the composition $F = f \circ H$ is a circuit of size $s + N^{O(1)}$, depth $d+3$, and bottom fan-in $O(\log N)$. Moreover, if the circuit for f is monotone with AND gates at the bottom level, then the bottom level of f can be merged with the top level of H to reduce the depth of $F = f \circ H$ to at most $(d+3) - 1 = d+2$. $\square$

We now obtain our lower bounds on the smooth threshold degree of $\mathbf{AC}^0$. We present two incomparable theorems here, both of which apply Lemma 5.55 in a recursive manner but with different base cases.

THEOREM 5.56. *Let $k \geq 0$ be a given integer. Then there is an (explicitly given) circuit family $\{f_n\}_{n=1}^{\infty}$, where $f_n\colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k$, bottom fan-in $O(\log n)$, and smooth threshold degree*

$$\deg_\pm\left(f_n, \exp\left(-c' \cdot \frac{n^{1-\frac{1}{k+1}}}{\log^{\frac{k(k-1)}{2(k+1)}} n}\right)\right) \geq c'' \cdot \frac{n^{1-\frac{1}{k+1}}}{\log^{\frac{k(k-1)}{2(k+1)}} n} \tag{5.5.160}$$

for some constants $c', c'' > 0$ and all $n \geq 2$.

Proof. The proof is by induction on k. The base case $k = 0$ corresponds to the family of "dictator" functions $x \mapsto x_1$, each of which has 1/2-smooth threshold degree 1 by Fact 2.8. For the inductive step, fix an explicit circuit family $\{f_n\}_{n=1}^{\infty}$ in which $f_n\colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k$, and smooth threshold degree (5.5.160) for some constants $c', c'' > 0$. Then taking $\alpha = \frac{1}{k+1}$ and $\beta = \frac{k(k-1)}{2(k+1)}$ in Lemma 5.55 produces an explicit circuit family $\{F_n\}_{n=1}^{\infty}$ in which $F_n\colon \{0,1\}^n \to \{0,1\}$

has polynomial size, depth $3k+3 = 3(k+1)$, and smooth threshold degree

$$\deg_{\pm}\left(F_n, \exp\left(-C' \cdot \frac{n^{\frac{k+1}{k+2}}}{\log^{\frac{k(k+1)}{2(k+2)}} n}\right)\right) \geq C'' \cdot \frac{n^{\frac{k+1}{k+2}}}{\log^{\frac{k(k+1)}{2(k+2)}} n}$$

for some constants $C', C'' > 0$. This completes the inductive step. □

Theorem 5.57. *Let $k \geq 1$ be a given integer. Then there is an* (*explicitly given*) *circuit family $\{f_n\}_{n=1}^{\infty}$, where $f_n\colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k+1$, bottom fan-in $O(\log n)$, and smooth threshold degree*

$$\deg_{\pm}\left(f_n, \exp\left(-c' \cdot \frac{n^{1-\frac{2}{2k+3}}}{\log^{\frac{k^2}{2k+3}} n}\right)\right) \geq c'' \cdot \frac{n^{1-\frac{2}{2k+3}}}{\log^{\frac{k^2}{2k+3}} n} \qquad (5.5.161)$$

for some constants $c', c'' > 0$ and all $n \geq 2$.

Proof. As with Theorem 5.56, the proof is by induction on k. For the base case $k=1$, consider the family $\{g_n\}_{n=1}^{\infty}$ in which $g_n\colon \{0,1\}^n \to \{0,1\}$ is given by

$$g_n(x) = \bigvee_{i=1}^{\lfloor n^{1/3} \rfloor} \bigwedge_{j=1}^{\lfloor n^{2/3} \rfloor} x_{i,j}.$$

Then

$$\begin{aligned} \deg_{\pm}(g_n, 12^{-\lfloor n^{1/3}\rfloor - 1}) &= \deg_{\pm}(\mathrm{MP}_{\lfloor n^{1/3}\rfloor, \lfloor n^{2/3}\rfloor}, 12^{-\lfloor n^{1/3}\rfloor - 1}) \\ &\geq cn^{1/3} \end{aligned}$$

for some absolute constant $c > 0$, where the first step is valid because a function's smooth threshold degree remains unchanged when one negates the function or its input variables, and the second step uses Theorem 5.36. Applying Lemma 5.55 to the circuit family $\{g_n\}_{n=1}^{\infty}$ with $\alpha = 2/3$ and $\beta = 0$ yields an explicit circuit family $\{G_n\}_{n=1}^{\infty}$ in which $G_n\colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $2+2=4$, bottom fan-in $O(\log n)$, and smooth threshold degree

$$\deg_\pm\left(G_n, \exp\left(-C' \cdot \frac{n^{3/5}}{\log^{1/5} n}\right)\right) \geq C'' \cdot \frac{n^{3/5}}{\log^{1/5} n}$$

for some constants $C', C'' > 0$. This new circuit family $\{G_n\}_{n=1}^\infty$ establishes the base case.

For the inductive step, fix an integer $k \geq 1$ and an explicit circuit family $\{f_n\}_{n=1}^\infty$ in which $f_n\colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k+1$, and smooth threshold degree (5.5.161) for some constants $c', c'' > 0$. Applying Lemma 5.55 with $\alpha = 2/(2k+3)$ and $\beta = k^2/(2k+3)$ yields an explicit circuit family $\{F_n\}_{n=1}^\infty$, where $F_n\colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $(3k+1)+3 = 3(k+1)+1$, bottom fan-in $O(\log n)$, and smooth threshold degree

$$\deg_\pm\left(F_n, \exp\left(-C''' \cdot \frac{n^{\frac{2k+3}{2k+5}}}{\log^{\frac{(k+1)^2}{2k+5}} n}\right)\right) \geq C'''' \cdot \frac{n^{\frac{2k+3}{2k+5}}}{\log^{\frac{(k+1)^2}{2k+5}} n}$$

for some constants $C''', C'''' > 0$. This completes the inductive step. □

5.5.8. The sign-rank of AC⁰. We have reached our main result on the sign-rank and unbounded-error communication complexity of constant-depth circuits. The proof amounts to lifting, by means of Theorem 3.11, the lower bounds on the smooth threshold degree in Theorems 5.56 and 5.57 to sign-rank lower bounds.

THEOREM 5.58. *Let $k \geq 1$ be a given integer. Then there is an (explicitly given) Boolean circuit family $\{F_n\}_{n=1}^\infty$, where $F_n\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k$, bottom fan-in $O(\log n)$, sign-rank*

$$\operatorname{rk}_\pm(F_n) = \exp\left(\Omega\left(n^{1-\frac{1}{k+1}} \cdot (\log n)^{-\frac{k(k-1)}{2(k+1)}}\right)\right), \tag{5.5.162}$$

and unbounded-error communication complexity

$$\mathsf{UPP}(F_n) = \Omega\left(n^{1-\frac{1}{k+1}} \cdot (\log n)^{-\frac{k(k-1)}{2(k+1)}}\right). \tag{5.5.163}$$

Proof. Theorem 5.56 constructs a circuit family $\{f_n\}_{n=1}^{\infty}$ in which $f_n\colon \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k$, bottom fan-in $O(\log n)$, and smooth threshold degree (5.5.160) for some constants $c', c'' > 0$ and all $n \geq 2$. Abbreviate $m = 2\lceil \exp(4c'/c'')\rceil$. For any $n \geq m$, define $F_n = f_{\lfloor n/m \rfloor} \circ \mathrm{OR}_m \circ \mathrm{AND}_2$. Then (5.5.162) is immediate from (5.5.160) and Theorem 3.11. Combining (5.5.163) with Theorem (3.9) settles (5.5.163).

It remains to analyze the circuit complexity of F_n. We defined F_n formally as a circuit of depth $3k+2$ in which the bottom four levels have fan-ins $n^{O(1)}$, $O(\log n)$, $2m$, and 2, in that order. Since m is a constant independent of n, these four levels can be computed by a circuit of polynomial size, depth 2, and bottom fan-in $O(\log n)$. This optimization reduces the depth of F_n to $(3k+2) - 4 + 2 = 3k$ while keeping the bottom fan-in at $O(\log n)$. □

We now similarly lift Theorem 5.57 to a lower bound on sign-rank and unbounded-error communication complexity.

THEOREM 5.59. *Let $k \geq 1$ be a given integer. Then there is an (explicitly given) Boolean circuit family $\{F_n\}_{n=1}^{\infty}$, where $F_n\colon \{0,1\}^n \times \{0,1\}^n \to \{0,1\}$ has polynomial size, depth $3k+1$, bottom fan-in $O(\log n)$, sign-rank*

$$\mathrm{rk}_{\pm}(F_n) = \exp\left(\Omega\left(n^{1-\frac{2}{2k+3}} \cdot (\log n)^{-\frac{k^2}{2k+3}}\right)\right),$$

and unbounded-error communication complexity

$$\mathsf{UPP}(F_n) = \Omega\left(n^{1-\frac{2}{2k+3}} \cdot (\log n)^{-\frac{k^2}{2k+3}}\right).$$

Proof. The proof is analogous to that of Theorem 5.58, with the only difference that the appeal to Theorem 5.56 should be replaced with an appeal to Theorem 5.57. □

Theorems 5.58 and 5.59 settle Theorems 5.2, 5.3, and 5.5 in the introduction.

5.6. A dual object for OR

The purpose of this section is to prove Theorem 5.17, which gives a dual polynomial for the OR function with a number of additional properties. The treatment here closely follows earlier work by Špalek [**137**], Bun and Thaler [**34, 38, 33**], and Sherstov [**122, 124**]. We start with a well-known binomial identity [**68**].

FACT 5.60. *For every univariate polynomial p of degree less than n,*

$$\sum_{t=0}^{n}(-1)^t\binom{n}{t}p(t)=0.$$

The next lemma constructs a dual polynomial for OR that has the sign behavior claimed in Theorem 5.17 but may lack some of the metric properties. The lemma is an adaptation of [**122**, Lemma A.2].

LEMMA 5.61. *Let ε be given, $0<\varepsilon<1$. Then for some constant $c=c(\varepsilon)\in(0,1)$ and every integer $n\geq 1$, there is an (explicitly given) function $\omega\colon\{0,1,2,\ldots,n\}\to\mathbb{R}$ such that*

$$\omega(0)>\frac{1-\varepsilon}{2}\cdot\|\omega\|_1, \tag{5.6.1}$$

$$|\omega(t)|\leq\frac{1}{ct^2\,2^{ct/\sqrt{n}}}\cdot\|\omega\|_1 \qquad (t=1,2,\ldots,n), \tag{5.6.2}$$

$$(-1)^t\omega(t)\geq 0 \qquad (t=0,1,2,\ldots,n), \tag{5.6.3}$$

$$\operatorname{orth}\omega\geq c\sqrt{n}. \tag{5.6.4}$$

Remark 5.62. It is helpful to keep in mind that properties (5.6.1)–(5.6.4) are logically monotonic in c. In other words, establishing these properties for a given constant $c > 0$ also establishes them for all smaller positive constants.

Proof of Lemma 5.61. Let $\Delta = 8\lceil 1/\varepsilon\rceil + 3$. If $n \le \Delta$, the requirements of the lemma hold for the function $\omega : (0, 1, 2, 3 \ldots, n) \mapsto (1, -1, 0, 0, \ldots, 0)$ and all $c \in (0, 1/\Delta]$. In what follows, we treat the complementary case $n > \Delta$.
Define $d = \lfloor\sqrt{n/\Delta}\rfloor$ and let $S = \{1, \frac{\Delta+1}{2}\} \cup \{i^2\Delta : i = 0, 1, 2, \ldots, d\}$, so that $S \subseteq \{0, 1, 2, \ldots, n\}$. Consider the function $\omega\colon \{0, 1, 2, \ldots, n\} \to \mathbb{R}$ given by

$$\omega(t) = \frac{(-1)^{n+t+|S|+1}}{n!}\binom{n}{t}\prod_{\substack{i=0,1,2,\ldots,n:\\ i\notin S}}(t-i).$$

Fact 5.60 implies that

$$\begin{aligned}\operatorname{orth}\omega &> d+1\\ &\ge \sqrt{\frac{n}{\Delta}}.\end{aligned}\tag{5.6.5}$$

A routine calculation reveals that

$$\omega(t) = \begin{cases}(-1)^{|\{i\in S: i<t\}|}\prod_{i\in S\setminus\{t\}}\frac{1}{|t-i|} & \text{if } t \in S,\\ 0 & \text{otherwise.}\end{cases}\tag{5.6.6}$$

It follows that

$$\begin{aligned}\frac{\omega(0)}{|\omega(1)|} &= \frac{\Delta-1}{\Delta+1}\prod_{i=1}^{d}\frac{i^2\Delta-1}{i^2\Delta} \\ &\geq 1-\frac{2}{\Delta+1}-\sum_{i=1}^{d}\frac{1}{i^2\Delta} \\ &> 1-\frac{2}{\Delta+1}-\frac{1}{\Delta}\sum_{i=1}^{\infty}\frac{1}{i^2} \\ &> 1-\frac{4}{\Delta}. \end{aligned} \tag{5.6.7}$$

An analogous application of (5.6.6) shows that

$$\begin{aligned}\frac{|\omega(\frac{\Delta+1}{2})|}{|\omega(0)|} &= \frac{\frac{\Delta+1}{2}}{\frac{\Delta+1}{2}\cdot(\frac{\Delta+1}{2}-1)}\frac{\Delta^d d!\, d!}{(\Delta-\frac{\Delta+1}{2})\cdot\frac{1}{2}\Delta^{d-1}(d-1)!\,(d+1)!} \\ &= \frac{8\Delta d}{(\Delta-1)^2(d+1)} \\ &\leq \frac{8\Delta}{(\Delta-1)^2}. \end{aligned} \tag{5.6.8}$$

Finally, for $i = 1, 2, \ldots, d$,

$$\begin{aligned}
\frac{|\omega(i^2\Delta)|}{|\omega(0)|} &= \frac{\frac{\Delta+1}{2}}{(i^2\Delta - 1)(i^2\Delta - \frac{\Delta+1}{2})} \cdot \frac{d!\, d!\, \Delta^d}{\frac{1}{2} \cdot (d-i)!\, (d+i)!\, \Delta^d} \\
&\leq \frac{2(\Delta+1)}{i^4(\Delta-1)^2} \cdot \frac{d!\, d!}{(d-i)!\, (d+i)!} \\
&= \frac{2(\Delta+1)}{i^4(\Delta-1)^2} \cdot \frac{d}{d+i} \cdot \frac{d-1}{d+i-1} \cdot \cdots \cdot \frac{d-i+1}{d+1} \\
&\leq \frac{2(\Delta+1)}{i^4(\Delta-1)^2} \cdot \left(1 - \frac{i}{d+i}\right)^i \\
&\leq \frac{2(\Delta+1)}{i^4(\Delta-1)^2} \cdot \exp\left(-\frac{i^2}{d+i}\right) \\
&\leq \frac{2(\Delta+1)}{i^4(\Delta-1)^2} \cdot \exp\left(-\frac{i^2}{2d}\right) \\
&\leq \frac{2(\Delta+1)}{i^4(\Delta-1)^2} \cdot \exp\left(-\frac{i^2}{2\sqrt{n/\Delta}}\right). \qquad (5.6.9)
\end{aligned}$$

Now,

$$\begin{aligned}
\frac{\|\omega\|_1}{\omega(0)} &= 1 + \frac{|\omega(1)|}{\omega(0)} + \frac{|\omega(\frac{\Delta+1}{2})|}{\omega(0)} + \sum_{i=1}^{d} \frac{|\omega(i^2\Delta)|}{\omega(0)} \\
&\leq 1 + \left(1 - \frac{4}{\Delta}\right)^{-1} + \frac{8\Delta}{(\Delta-1)^2} + \sum_{i=1}^{\infty} \frac{2(\Delta+1)}{i^4(\Delta-1)^2} \\
&= 1 + \left(1 - \frac{4}{\Delta}\right)^{-1} + \frac{8\Delta}{(\Delta-1)^2} + \frac{\pi^4(\Delta+1)}{45(\Delta-1)^2} \\
&\leq \frac{2}{1 - \frac{8}{\Delta}} \\
&< \frac{2}{1-\varepsilon}, \qquad (5.6.10)
\end{aligned}$$

where the second step uses (5.6.7)–(5.6.9), and the last step substitutes the definition of Δ. Now (5.6.1) follows from (5.6.10). Moreover, for $c = c(\Delta) > 0$ small enough, (5.6.4) follows from (5.6.5), whereas (5.6.2) follows from (5.6.9) and the fact that ω vanishes outside the union $\{1, \frac{\Delta+1}{2}\} \cup \{i^2\Delta : i = 0, 1, 2, \ldots, d\}$.

It remains to verify that ω has the desired sign behavior. Since ω vanishes outside S, the requirement (5.6.3) holds trivially at those points. For $t \in S$, it follows from (5.6.6) that

$$\begin{aligned}
&\operatorname{sgn} \omega(1) = -1, \\
&\operatorname{sgn} \omega\big(\tfrac{\Delta+1}{2}\big) = 1, \\
&\operatorname{sgn} \omega(i^2\Delta) = (-1)^i, && i = 0, 1, 2, \ldots, d.
\end{aligned}$$

Since $\Delta \in 4\mathbb{Z} + 3$ by definition, we conclude that $\operatorname{sgn} \omega(t) = (-1)^t$ for all $t \in S$. This settles (5.6.3) and completes the proof. □

We have reached the main result of this section.

THEOREM (restatement of Theorem 5.17). *Let $0 < \varepsilon < 1$ be given. Then for some constants $c', c'' \in (0, 1)$ and all integers $N \geq n \geq 1$, there is an (explicitly given) function $\psi \colon \{0, 1, 2, \ldots, N\} \to \mathbb{R}$ such that*

$$\psi(0) > \frac{1-\varepsilon}{2}, \tag{5.6.11}$$

$$\|\psi\|_1 = 1, \tag{5.6.12}$$

$$\operatorname{orth} \psi \geq c'\sqrt{n}, \tag{5.6.13}$$

$$\operatorname{sgn} \psi(t) = (-1)^t, \qquad t = 0, 1, 2, \ldots, N, \tag{5.6.14}$$

$$|\psi(t)| \in \left[\frac{c'}{(t+1)^2\, 2^{c''t/\sqrt{n}}}, \; \frac{1}{c'(t+1)^2\, 2^{c''t/\sqrt{n}}}\right], \qquad t = 0, 1, 2, \ldots, N. \tag{5.6.15}$$

Proof. The degenerate case $N = 1$ holds for the function $\omega : (0, 1) \mapsto (1/2, -1/2)$ and all $c', c'' \in (0, 1/4)$. In the rest of the proof, we treat the complementary case $N \geq 2$.

For some sufficiently small constant $c \in (0, 1/4)$ and all $n \geq 1$, Lemma 5.61 and Remark 5.62 ensure the existence of a function $\omega\colon \{0, 1, 2, \ldots, \lceil n/2 \rceil\} \to \mathbb{R}$ such that

$$\|\omega\|_1 = 1, \tag{5.6.16}$$

$$\omega(0) > \frac{1}{2}\left(1 - \frac{\varepsilon}{6}\right), \tag{5.6.17}$$

$$|\omega(t)| \leq \frac{1}{ct^2\, 2^{ct/\sqrt{n}}} \qquad (t = 1, 2, \ldots, \lceil n/2 \rceil), \tag{5.6.18}$$

$$(-1)^t \omega(t) \geq 0 \qquad (t = 0, 1, 2, \ldots, \lceil n/2 \rceil), \tag{5.6.19}$$

$$\operatorname{orth} \omega \geq c\sqrt{n}. \tag{5.6.20}$$

For convenience, extend ω to all of $\mathbb{Z}$ by defining it to be zero outside its original domain. Define $\Psi\colon \{0, 1, 2, \ldots, N\} \to \mathbb{R}$ by

$$\Psi(t) = \omega(t) + \delta \left(\sum_{i=1}^{N - \lceil n/2 \rceil} \frac{(-1)^i}{i^2\, 2^{ci/\sqrt{n}}} \omega(t - i) \right. \\ \left. + \sum_{i = N - \lceil n/2 \rceil + 1}^{N} \frac{(-1)^i}{i^2\, 2^{ci/\sqrt{n}}} \omega(-t + i) \right),$$

where

$$\delta = \frac{5\varepsilon}{\pi^2 (1 - \varepsilon)}.$$

By (5.6.20) and Proposition 5.7(i),

$$\operatorname{orth} \Psi \geq c\sqrt{n}. \tag{5.6.21}$$

We now move on to metric properties of Ψ. Multiplying the defining equation for Ψ on both sides by $(-1)^t$ and applying (5.6.19), we arrive at

$$(-1)^t\Psi(t) = |\omega(t)| + \delta \left(\sum_{i=1}^{N-\lceil n/2\rceil} \frac{|\omega(t-i)|}{i^2\, 2^{ci/\sqrt{n}}} + \sum_{i=N-\lceil n/2\rceil+1}^{N} \frac{|\omega(-t+i)|}{i^2\, 2^{ci/\sqrt{n}}} \right), \qquad t = 0,1,2,\ldots,N. \quad (5.6.22)$$

Summing over t gives

$$\begin{aligned} \|\Psi\|_1 &= \|\omega\|_1 + \delta \sum_{i=1}^{N} \frac{1}{i^2\, 2^{ci/\sqrt{n}}} \|\omega\|_1 \\ &= 1 + \delta \sum_{i=1}^{N} \frac{1}{i^2\, 2^{ci/\sqrt{n}}} \\ &\in \left[1, 1 + \delta \sum_{i=1}^{\infty} \frac{1}{i^2}\right] \\ &= \left[1, \frac{6-\varepsilon}{6(1-\varepsilon)}\right], \end{aligned} \quad (5.6.23)$$

where the second step uses (5.6.16). We also have

$$\begin{aligned} \Psi(0) &\geq \omega(0) \\ &> \frac{6-\varepsilon}{12}, \end{aligned} \quad (5.6.24)$$

where the first and second steps use (5.6.22) and (5.6.17), respectively.

We now estimate $|\Psi(t)|$ for each $t = 1, 2, \ldots, N$. For a lower bound, we have

$$\begin{aligned} |\Psi(t)| &= |\omega(t)| + \delta \left(\sum_{i=1}^{N-\lceil n/2 \rceil} \frac{|\omega(t-i)|}{i^2\, 2^{ci/\sqrt{n}}} + \sum_{i=N-\lceil n/2 \rceil+1}^{N} \frac{|\omega(-t+i)|}{i^2\, 2^{ci/\sqrt{n}}} \right) \\ &\geq \delta \cdot \frac{|\omega(0)|}{t^2\, 2^{ct/\sqrt{n}}} \\ &\geq \frac{5\varepsilon}{\pi^2(1-\varepsilon)} \cdot \frac{6-\varepsilon}{12} \cdot \frac{1}{t^2\, 2^{ct/\sqrt{n}}}, \end{aligned} \tag{5.6.25}$$

where the first and last steps use (5.6.22) and (5.6.17), respectively. The upper bound on $|\Psi(t)|$ is somewhat more technical. To begin with, we have the following bound for every positive integer t:

$$\begin{aligned} \sum_{i=1}^{t-1} \frac{1}{(t-i)^2\, i^2} &= \sum_{i=1}^{t-1} \frac{1}{\max\{(t-i)^2, i^2\}\ \min\{(t-i)^2, i^2\}} \\ &\leq \frac{1}{(t/2)^2} \sum_{i=1}^{t-1} \frac{1}{\min\{(t-i)^2, i^2\}} \\ &\leq \frac{1}{(t/2)^2} \cdot 2 \sum_{i=1}^{\infty} \frac{1}{i^2} \\ &= \frac{4\pi^2}{3t^2}. \end{aligned} \tag{5.6.26}$$

Continuing,

$$\begin{aligned} \sum_{i=1}^{\infty} \frac{|\omega(t-i)|}{i^2\, 2^{ci/\sqrt{n}}} &= \frac{|\omega(0)|}{t^2\, 2^{ct/\sqrt{n}}} + \sum_{i=1}^{t-1} \frac{|\omega(t-i)|}{i^2\, 2^{ci/\sqrt{n}}} \\ &\leq \frac{1}{t^2\, 2^{ct/\sqrt{n}}} + \sum_{i=1}^{t-1} \frac{1}{c(t-i)^2\, i^2\, 2^{ct/\sqrt{n}}} \\ &\leq \frac{1}{t^2\, 2^{ct/\sqrt{n}}} \left(1 + \frac{4\pi^2}{3c} \right), \end{aligned} \tag{5.6.27}$$

where the second step uses (5.6.16) and (5.6.18), and the third step substitutes the bound from (5.6.26). Analogously,

$$\begin{aligned}\sum_{i=1}^{\infty} \frac{|\omega(-t+i)|}{i^2\, 2^{ci/\sqrt{n}}} &= \frac{|\omega(0)|}{t^2\, 2^{ct/\sqrt{n}}} + \sum_{i=t+1}^{\infty} \frac{|\omega(-t+i)|}{i^2\, 2^{ci/\sqrt{n}}} \\ &\leq \frac{1}{t^2\, 2^{ct/\sqrt{n}}} + \sum_{i=t+1}^{\infty} \frac{1}{c(t-i)^2\, i^2\, 2^{ci/\sqrt{n}}} \\ &\leq \frac{1}{t^2\, 2^{ct/\sqrt{n}}} \left(1 + \sum_{i=t+1}^{\infty} \frac{1}{c(t-i)^2}\right) \\ &\leq \frac{1}{t^2\, 2^{ct/\sqrt{n}}} \left(1 + \frac{\pi^2}{6c}\right), \end{aligned} \tag{5.6.28}$$

where the second step uses (5.6.16) and (5.6.18). Now for every integer $t \geq 1$,

$$\begin{aligned}|\Psi(t)| &\leq |\omega(t)| + \delta \left(\sum_{i=1}^{\infty} \frac{|\omega(t-i)|}{i^2\, 2^{ci/\sqrt{n}}} + \sum_{i=1}^{\infty} \frac{|\omega(-t+i)|}{i^2\, 2^{ci/\sqrt{n}}} \right) \\ &\leq \frac{1}{ct^2\, 2^{ct/\sqrt{n}}} \left(1 + 2c\delta + \frac{4\pi^2\delta}{3} + \frac{\pi^2\delta}{6}\right), \end{aligned} \tag{5.6.29}$$

where the first step is immediate from the defining equation for Ψ, and the second step uses (5.6.18), (5.6.27), and (5.6.28). To complete the proof, let $\psi\colon \{0,1,2,\ldots,N\} \to \mathbb{R}$ be given by $\psi = \Psi/\|\Psi\|_1$. Then for a small enough constant $c' = c'(c,\varepsilon,\delta) > 0$ and $c'' = c$, properties (5.6.11)–(5.6.15) follow directly from (5.6.21)–(5.6.25) and (5.6.29). □

5.7. Sign-rank and smooth threshold degree

The purpose of this section is to prove Theorem 3.11, implicit in [**117, 106**]. We closely follow the treatment in those earlier papers. Sections 5.7.1–5.7.2 cover necessary technical background, followed by the proof proper in Section 5.7.3.

5.7.1. Forster's bound. The *spectral norm* of a real matrix $A = [A_{xy}]_{x\in X, y\in Y}$ is given by

$$\|A\| = \max_{v\in\mathbb{R}^{|Y|},\ \|v\|_2=1} \|Av\|_2,$$

where $\|\cdot\|_2$ is the Euclidean norm on vectors. The first strong lower bound on the sign-rank of an explicit matrix was obtained by Forster [**53**], who proved that

$$\mathrm{rk}_\pm(A) \geq \frac{\sqrt{|X|\,|Y|}}{\|A\|}$$

for any matrix $A = [A_{xy}]_{x\in X, y\in Y}$ with ±1 entries. Forster's result has seen a number of generalizations, including the following theorem due to Forster et al. [**54**, Theorem 3].

THEOREM 5.63 (Forster et al.). *Let $A = [A_{xy}]_{x\in X, y\in Y}$ be a real matrix without zero entries. Then*

$$\mathrm{rk}_\pm(A) \geq \frac{\sqrt{|X|\,|Y|}}{\|A\|}\ \min_{x,y} |A_{xy}|.$$

5.7.2. Spectral norm of pattern matrices. *Pattern matrices* were introduced in [**114, 116**] and proved useful in obtaining strong lower bounds on communication complexity. Relevant definitions and results from [**116**] follow. Let n and N be positive integers with $n \mid N$. Partition $\{1,2,\dots,N\}$ into n contiguous blocks, each with N/n elements:

$$\{1,2,\dots,N\} = \left\{1,2,\dots,\frac{N}{n}\right\} \cup \left\{\frac{N}{n}+1,\dots,\frac{2N}{n}\right\} \\ \cup\cdots\cup \left\{\frac{(n-1)N}{n}+1,\dots,N\right\}.$$

Now, let $\mathcal{V}(N,n)$ denote the family of subsets $V \subseteq \{1,2,\dots,N\}$ that have exactly one element in each of these blocks (in particular, $|V| = n$). Clearly, $|\mathcal{V}(N,n)| = (N/n)^n$. For a function $\phi\colon \{0,1\}^n \to \mathbb{R}$, the (N,n,ϕ)-*pattern matrix* is the real matrix A given by

$$A = \Big[\phi(x|_V \oplus w)\Big]_{x\in\{0,1\}^N,\,(V,w)\in\mathcal{V}(N,n)\times\{0,1\}^n}.$$

In words, A is the matrix of size 2^N by $(N/n)^n 2^n$ whose rows are indexed by strings $x \in \{0,1\}^N$, whose columns are indexed by pairs $(V,w) \in \mathcal{V}(N,n)\times\{0,1\}^n$, and whose entries are given by $A_{x,(V,w)} = \phi(x|_V \oplus w)$. We will need the following expression for the spectral norm of a pattern matrix [**116**, Theorem 4.3].

THEOREM 5.64 (Sherstov). *Let $\phi\colon \{0,1\}^n \to \mathbb{R}$ be given. Let A be the (N,n,ϕ)-pattern matrix. Then*

$$\|A\| \;=\; \sqrt{2^{N+n}\left(\frac{N}{n}\right)^n}\;\max_{S\subseteq\{1,2,\ldots,n\}}\left\{|\hat\phi(S)|\left(\frac{n}{N}\right)^{|S|/2}\right\}.$$

5.7.3. Proof of Theorem 3.11. We are now in a position to prove Theorem 3.11. We will derive it from the following more general result, stated in terms of pattern matrices.

THEOREM 5.65. *Let $f\colon \{0,1\}^n \to \{0,1\}$ be given. Suppose that $\deg_\pm(f,\gamma) \geq d$, where γ and d are positive reals. Then for any integer $T \geq 1$, the $(Tn, n, (-1)^f)$-pattern matrix has sign-rank at least $\gamma T^{d/2}$.*

Proof. By the definition of smooth threshold degree, there is a probability distribution μ on $\{0,1\}^n$ such that

$$\mu(x) \geq \gamma\, 2^{-n}, \qquad x \in \{0,1\}^n, \tag{5.7.1}$$

$$\operatorname{orth}((-1)^f \cdot \mu) \geq d. \tag{5.7.2}$$

Abbreviate $\phi = (-1)^f \cdot \mu$. Let F and Φ denote the $(Tn, n, (-1)^f)$- and (Tn, n, ϕ)-pattern matrices, respectively. By (5.2.3) and (5.7.2),

$$\hat{\phi}(S) = 0, \qquad\qquad |S| < d. \tag{5.7.3}$$

The remaining Fourier coefficients of ϕ can be bounded using Proposition 2.1:

$$|\hat{\phi}(S)| \le 2^{-n}, \qquad\qquad S \subseteq \{1, 2, \ldots, n\}. \tag{5.7.4}$$

Now

$$\begin{aligned}\mathrm{rk}_{\pm}(F) &= \mathrm{rk}_{\pm}(\Phi)\\ &\ge \frac{\sqrt{2^{Tn+n}\, T^n}}{\|\Phi\|} \cdot \gamma\, 2^{-n}\\ &= \frac{\gamma\, 2^{-n}}{\max_S\{|\hat{\phi}(S)|\, T^{-|S|/2}\}}\\ &\ge \gamma T^{d/2},\end{aligned}$$

where the first step is valid because F and Φ have the same sign pattern; the second step uses (5.7.1) and Theorem 5.63; the third step applies Theorem 5.64; and the final step substitutes the upper bounds from (5.7.3) and (5.7.4). □

We have reached the main result of this section.

THEOREM (restatement of Theorem 3.11). *Let $f\colon \{0,1\}^n \to \{0,1\}$ be given. Suppose that* $\deg_{\pm}(f, \gamma) \ge d$*, where γ and d are positive reals. Fix an integer $m \ge 2$ and define $F\colon \{0,1\}^{mn} \times \{0,1\}^{mn} \to \{0,1\}$ by $F(x, y) = f \circ \mathrm{OR}_m \circ \mathrm{AND}_2$. Then*

$$\mathrm{rk}_{\pm}(F) \ge \gamma \left\lfloor \frac{m}{2} \right\rfloor^{d/2}.$$

Proof. The result is immediate from Theorem 5.65 since the $(\lfloor m/2 \rfloor n, n, (-1)^f)$-pattern matrix is a submatrix of $[(-1)^{F(x,y)}]_{x,y}$. □

CHAPTER 6

Randomized and quantum communication complexity

In this chapter, we discuss our near-optimal separation of randomized and quantum communication complexity. Our approach is to first prove an optimal separation in the query model. Then we use the standard machinery to "lift" the separation from the query model to the communication model.

6.1. Introduction

Understanding the relative power of quantum and classical computing is of basic importance in theoretical computer science. This question has been studied most actively in the *query model*, which is tractable enough to allow unconditional lower bounds yet rich enough to capture most of the known quantum algorithms. Illustrative examples include the quantum algorithms of Deutsch and Jozsa [**50**], Bernstein and Vazirani [**18**], Grover [**69**], and Shor's period-finding [**130**]. In the query model, the task is to evaluate a fixed function f on an unknown n-bit input x. In the classical setting, query algorithms are commonly referred to as *decision trees.* A decision tree accesses the input one bit at a time, choosing the bits to query in adaptive fashion. The objective is to determine $f(x)$ by querying as few bits as possible. The minimum number of queries needed to determine $f(x)$ in the worst case is called the *query complexity of f*. The quantum model is a far-reaching generalization of the classical decision tree whereby all bits can be queried in superposition with a single query. The catch is that the outcomes of those queries are then also in superposition, and it is not clear a priori whether quantum query algorithms are more powerful than decision trees. The focus of our work is on the *bounded-error* regime, where the query algorithm (quantum or classical) is allowed to err with small constant probability on any given input.

The comparative power of randomized and quantum query algorithms has been studied for more than two decades. In pioneering work, Deutsch and Jozsa [**50**] gave a quantum query algorithm that solves, with a single query, a problem on n bits that any deterministic decision tree needs at least $n/2$ queries to solve. Unfortunately, this separation does not apply to the more subtle, bounded-error setting. This was addressed in follow-up work by Simon [**131**], who exhibited a problem with bounded-error quantum query complexity $O(\log^2 n)$ and randomized query complexity $\Omega(\sqrt{n})$. These are striking examples of the computational advantages afforded by the quantum model.

6.1.1. Forrelation and rorrelation. The above results leave us with a fundamental question: what is the largest possible separation between bounded-error quantum and randomized query complexity, for a problem with n-bit input? This question was raised by Buhrman et al. [**30**] and, a decade later, by Aaronson and Ambainis [**2**], who presented it as being essential to understanding the phenomenon of quantum speedups. Toward this goal, the authors of [**2**] obtained both positive and negative results. They showed, for every constant t, that every quantum algorithm with t queries can be converted to a randomized decision tree of cost $O(n^{1-1/2t})$. In particular, this rules out an $O(1)$ versus $\Omega(n)$ separation. In the opposite direction, Aaronson and Ambainis exhibited a problem that can be solved to bounded error with a single quantum query but has randomized query complexity $\tilde{\Omega}(\sqrt{n})$. They left open the challenge of obtaining a separation of $O(1)$ versus $\Omega(n^{\alpha})$ for some $\alpha > 1/2$.

In more detail, Aaronson and Ambainis [**2**] introduced and studied the *k-fold forrelation problem.* The input to the problem is a k-tuple of vectors $x_1, x_2, \ldots, x_k \in \{-1,1\}^n$, where n is a power of 2. Define

$$\phi_{n,k}(x_1, x_2, \ldots, x_k) = \frac{1}{n}\mathbf{1}^\mathsf{T} D_{x_1} H D_{x_2} H D_{x_3} H \cdots H D_{x_k} \mathbf{1}, \tag{6.1.1}$$

where $\mathbf{1}$ is the all-ones vector, H is the Hadamard transform matrix of order n, and D_{x_i} is the diagonal matrix with the vector x_i on the diagonal. Since each of the linear transformations $H, D_{x_1}, D_{x_2}, \ldots, D_{x_n}$ preserves Euclidean length, it follows that $|\phi_{n,k}(x_1, x_2, \ldots, x_k)| \leq 1$. Given $x_1, x_2, \ldots, x_k$, the forrelation problem is to distinguish between the cases $|\phi_{n,k}(x_1, x_2, \ldots, x_k)| \leq \alpha$ and $\phi_{n,k}(x_1, x_2, \ldots, x_k) \geq \beta$, where the problem parameters $0 < \alpha < \beta < 1$ are suitably chosen constants. Equation (6.1.1) directly gives a quantum algorithm that solves the forrelation problem with bounded error and query cost k, where the k queries correspond to the k diagonal matrices. The cost can be further reduced to $\lceil k/2 \rceil$ by viewing (6.1.1) as the *inner product* of two vectors obtained by $\lceil k/2 \rceil$ and $\lfloor k/2 \rfloor$ applications, respectively, of diagonal matrices [**2**]. Aaronson and Ambainis complemented this with an $\tilde{\Omega}(\sqrt{n})$ lower bound on the randomized query complexity of the forrelation problem for $k = 2$, hence the 1 versus $\tilde{\Omega}(\sqrt{n})$ separation mentioned above.

Building on the work of Aaronson and Ambainis [**2**], last year Tal [**135**] gave an improved separation of $O(1)$ versus $\Omega(n^{2/3-\varepsilon})$ for bounded-error quantum and randomized query complexities, for any constant $\varepsilon > 0$. For this, Tal replaced (6.1.1) with the more general quantity

$$\phi_{n,k,U}(x_1, x_2, \ldots, x_k) = \frac{1}{n}\mathbf{1}^\mathsf{T} D_{x_1} U D_{x_2} U D_{x_3} U \cdots U D_{x_k} \mathbf{1}, \tag{6.1.2}$$

where U is an arbitrary but fixed orthogonal matrix. On input $x_1, x_2, \ldots, x_k \in \{-1,1\}^n$, the author of [**135**] considered the problem of distinguishing between the cases $|\phi_{n,k,U}(x_1, x_2, \dots, x_k)| \le 2^{-k-1}$ and $\phi_{n,k,U}(x_1, x_2, \dots, x_k) \ge 2^{-k}$. This problem is referred to in [**135**] as the *k-fold rorrelation problem with respect to U*. The quantum algorithm of Aaronson and Ambainis, adapted to the arbitrary choice of U, solves this new problem with $\lceil k/2 \rceil$ queries and advantage $\Omega(2^{-k})$ over random guessing, which counts as a bounded-error algorithm for any constant k. On the other hand, Tal [**135**] proved that the randomized query complexity of the k-fold rorrelation problem for uniformly random U is $\Omega(n^{2(k-1)/(3k-1)}/k \log n)$ with high probability. While this is weaker than Aaronson and Ambainis's bound for $k = 2$, setting k to a large constant gives a separation of $O(1)$ versus $\Omega(n^{2/3-\varepsilon})$ for bounded-error quantum and randomized query complexity for any constant $\varepsilon > 0$.

6.1.2. Our results. Prior to our work, Tal's separation of $O(1)$ versus $\Omega(n^{2/3-\varepsilon})$ was the strongest known, and Aaronson and Ambainis's challenge of obtaining an $O(1)$ versus $\Omega(n^{1-\varepsilon})$ separation remained open. The main contribution of our work is to resolve this question.

Separations for partial functions. In what follows, we let $f_{n,k,U}$ denote the k-fold rorrelation problem with respect to U. We prove:

THEOREM 6.1. *Let n and k be positive integers, with $k \le \frac{1}{3}\log n - 1$. Let $U \in \mathbb{R}^{n\times n}$ be a uniformly random orthogonal matrix. Then with probability* $1 - o(1)$,

$$R^{\mathrm{dt}}_{\frac{1}{2}-\gamma}(f_{n,k,U}) = \Omega\left(\frac{\gamma^2}{k} \cdot \frac{n^{1-\frac{1}{k}}}{(\log n)^{2-\frac{1}{k}}}\right) \tag{6.1.3}$$

for all $0 \le \gamma \le 1/2$.

For $k = 2$, this lower bound is the same as Aaronson and Ambainis's lower bound for the forrelation problem (which is $f_{n,2,H}$ in our notation). For $k = 3$ already, Theorem 6.1 is a polynomial improvement on all previous work, including Tal's recent result [**135**]. Theorem 6.1 is essentially tight for all k, both even and odd, due to the matching upper bound $O_k(n^{1-1/k})$ [**2, 27**] . Since $f_{n,k,U}$ has an efficient quantum protocol for every U (see Section 6.5.2 for details), we obtain the following corollary:

Corollary 6.2. *Let $\varepsilon > 0$ be given. Then there is a partial Boolean function f on $\{-1,1\}^n$ with*

$$Q^{\mathrm{dt}}_{1/3}(f) = O(1),$$
$$R^{\mathrm{dt}}_{1/3}(f) = \Omega(n^{1-\varepsilon}).$$

This separation of bounded-error quantum and randomized query complexities is best possible for all f due to the aforementioned result that every quantum protocol with k queries can be simulated by a randomized query algorithm of cost $O(n^{1-1/2k})$. In particular, Corollary 6.2 shows that the rorrelation problem separates quantum and randomized query complexity optimally, of all problems f. The following incomparable corollary can be obtained by taking $k = k(n)$ in Theorem 6.1 to be an arbitrarily slow-growing function, e.g., $k = \log\log\log n$:

Corollary 6.3. *Let $\alpha\colon \mathbb{N} \to \mathbb{N}$ be any monotone function with $\alpha(n) \to \infty$ as $n \to \infty$. Then there is a partial Boolean function f on $\{-1,1\}^n$ with*

$$Q^{\mathrm{dt}}_{1/3}(f) \leq \alpha(n),$$
$$R^{\mathrm{dt}}_{1/3}(f) \geq n^{1-o(1)}.$$

Again, this quantum-classical separation is best possible since [**2, 27**] rules out the possibility of an $O(1)$ versus $n^{1-o(1)}$ gap.

A satisfying probability-theoretic interpretation of our results is that the phenomenon of quantum-classical gaps is a common one. More precisely, our results show that the set of orthogonal matrices U for which $f_{n,k,U}$ does *not* exhibit a best-possible quantum-classical separation has Haar measure 0. Prior to our work, this was unknown for any integer $k > 2$.

Separation for total functions. Our results so far pertain to *partial* Boolean functions, whose domain of definition is a proper subset of the Boolean hypercube. For total Boolean functions, such large quantum-classical gaps are not possible. In a seminal paper, Beals et al. [**13**] prove that the bounded-error quantum query complexity of a total function f is always polynomially related to the randomized query complexity of f. A natural question to ask is how large this polynomial gap can be. Grover's search [**69**] shows that the n-bit OR function has bounded-error quantum query complexity $\Theta(\sqrt{n})$ and randomized complexity $\Theta(n)$. For a long time, this quadratic separation was believed to be the largest possible. In a surprising result, Aaronson et al. [**3**] proved the existence of a total function f with $R^{\mathrm{dt}}_{1/3}(f) = \tilde{\Omega}(Q^{\mathrm{dt}}_{1/3}(f)^{2.5})$. This was improved by Tal [**135**] to $R^{\mathrm{dt}}_{1/3}(f) \geq Q^{\mathrm{dt}}_{1/3}(f)^{8/3-o(1)}$. We give a polynomially stronger separation:

Theorem 6.4. *There is a function $f\colon \{-1,1\}^n \to \{0,1\}$ with*

$$R^{\mathrm{dt}}_{1/3}(f) \geq Q^{\mathrm{dt}}_{1/3}(f)^{3-o(1)}.$$

Theorem 6.4 follows automatically by combining our Corollary 6.3 with the "cheat-sheet" framework of Aaronson et al. [**3**]. Specifically, they prove that any partial function f on n bits that exhibits an $n^{o(1)}$ versus $n^{1-o(1)}$ separation for bounded-error quantum versus randomized query complexity, can be automatically converted into a total function with $R^{\mathrm{dt}}_{1/3}(f) \geq Q^{\mathrm{dt}}_{1/3}(f)^{3-o(1)}$. A recent paper of Aaronson et al. [**4**] conjectures that $R^{\mathrm{dt}}_{1/3}(f) = O(Q^{\mathrm{dt}}_{1/3}(f)^3)$ for every total function f, which would mean that our separation in Theorem 6.4 is essentially optimal. The best current upper bound is $R^{\mathrm{dt}}_{1/3}(f) = O(Q^{\mathrm{dt}}_{1/3}(f)^4)$ due to [**4**], derived there from the breakthrough result of Huang [**72**] on the sensitivity conjecture.

Separations for communication complexity. Using standard reductions, our quantum-classical query separations imply analogous separations for communication complexity. In more detail, let f be a (possibly partial) Boolean function on $\{-1,1\}^n$. For any communication problem $g\colon \{-1,1\}^m \times \{-1,1\}^m \to \{-1,1\}$, we let $f \circ g$ denote the (possibly partial) communication problem on $(\{-1,1\}^m)^n \times (\{-1,1\}^m)^n$ given by $(f \circ g)(x,y) = f(g(x_1,y_1), g(x_2,y_2), \ldots, g(x_n,y_n))$. Buhrman, Cleve, and Wigderson [**28**] proved that any quantum query algorithm for f gives a quantum communication protocol for $f \circ g$ with the same error and approximately the same cost. Quantitatively,

$$Q_\varepsilon(f \circ g) \le Q^{\mathrm{dt}}_\varepsilon(f) \cdot O(m + \log n), \tag{6.1.4}$$

where Q_ε denotes ε-error quantum communication complexity. Reversing this inequality has seen a great deal of work, mainly in the classical setting. In light of query-to-communication lifting, our main results have the following consequences.

Theorem 6.5. *Let $\varepsilon > 0$ be given. Then there is a partial Boolean function F on $\{-1,1\}^N \times \{-1,1\}^N$ with*

$$Q_{1/3}(F) = O(\log N),$$
$$R_{1/3}(F) = \Omega(N^{1-\varepsilon}).$$

Proof. Take f as in Corollary 6.2 and define $N = cn\log n$ and $F = f \circ \mathrm{IP}_{c\log n}$. Then the communication bounds follow from (6.1.4) and (3.2.1), respectively. □

Theorem 6.5 is essentially optimal and a polynomial improvement on previous work. The best previous quantum-classical separation for communication complexity was

$O(\log N)$ versus $\Omega(N^{2/3-\varepsilon})$, implicit in Tal [**135**] and preceded in turn by other exponential separations [**103, 107, 56**]. Similarly, our Corollary 6.3 translates in a black-box manner to communication complexity:

THEOREM 6.6. *Let* $\alpha\colon \mathbb{N} \to \mathbb{N}$ *be any monotone function with* $\alpha = \omega(1)$. *Then there is a partial Boolean function* F *on* $\{-1,1\}^N \times \{-1,1\}^N$ *with*

$$Q_{1/3}(F) \leq \alpha(N) \log N,$$

$$R_{1/3}(F) \geq N^{1-o(1)}.$$

Proof. Take f as in Corollary 6.3 and define $N = cn \log n$ and $F = f \circ \mathrm{IP}_{c \log n}$. Then the communication bounds follow from (6.1.4) and (3.2.1), respectively. □

Finally, we obtain the following result for *total* functions.

THEOREM 6.7. *There is a function* $F\colon \{-1,1\}^N \times \{-1,1\}^N \to \{0,1\}$ *with*

$$R_{1/3}(F) \geq Q_{1/3}(F)^{3-o(1)}.$$

Proof. The cheatsheet framework [**3**] ensures that the quantum and classical query complexities of f in Theorem 6.4 are polynomial in the number of variables n. With this in mind, we proceed as before, setting $N = cn \log n$ and $F = f \circ \mathrm{IP}_{c \log n}$ and applying (6.1.4) and (3.2.1). □

Again, Theorem 6.7 is a polynomial improvement on previous work, the best previous result being a power of 8/3 separation implicit in [**135**].

Fourier weight of decision trees. It is straightforward to verify that a uniformly random input $x \in (\{-1,1\}^n)^k$ is with high probability a *negative* instance of the rorrelation problem $f_{n,k,U}$. With this in mind, Tal [**135**] proves his lower bound for rorrelation by constructing a probability distribution $\mathcal{D}_{n,k,U}$ that generates *positive* instances of $f_{n,k,U}$ with nontrivial probability yet is indistinguishable from the uniform distribution by a decision tree T of cost $n^{2/3-O(1/k)}$. His notion of indistinguishability is based on the Fourier spectrum. Specifically, Tal [**135**] shows that: (i) the *sum* of the absolute values of the Fourier coefficients of T of given order ℓ does not grow too fast with ℓ; and (ii) the *maximum* Fourier coefficient of $\mathcal{D}_{n,k,U}$ of order ℓ decays exponentially fast with ℓ. In Tal's paper, the bound for (ii) is essentially optimal, whereas the bound for (i) is far from tight. The sum of the absolute values of the order-ℓ Fourier coefficients of a decision tree T, which we refer to as the *ℓ-Fourier weight of T*, is shown in [**135**] to be at most

$$c^\ell \sqrt{d^\ell (1 + \log kn)^{\ell-1}}, \tag{6.1.5}$$

where d is the depth of the tree and $c \geq 1$ is an absolute constant. This bound is strong for any constant ℓ but degrades rapidly as ℓ grows. In particular, for $\ell = \sqrt{d}$ already, (6.1.5) is weaker than the trivial bound $\binom{d}{\ell}$. This is a major obstacle since the indistinguishability proof requires strong bounds for every ℓ. This obstacle is the reason why Tal's analysis gives the randomized query lower bound $n^{2/3-O(1/k)}$ as opposed to the optimal $\tilde{\Omega}(n^{1-1/k})$. Tal conjectured that the ℓ-Fourier weight of a depth-d decision tree is in fact bounded by $c^\ell \sqrt{\binom{d}{\ell}(1 + \log kn)^{\ell-1}}$, which is a factor of $\sqrt{\ell!}$ improvement on (6.1.5) and essentially optimal. We prove his conjecture:

THEOREM 6.8. *Let $T\colon \{-1,1\}^n \to \{0,1\}$ be a function computable by a decision tree of depth d. Then*

$$\sum_{\substack{S \subseteq \{1,2,\ldots,n\}:\\ |S|=\ell}} |\hat{T}(S)| \leq c^\ell \sqrt{\binom{d}{\ell}(1 + \log n)^{\ell-1}}, \qquad \ell = 1, 2, \ldots, n,$$

where $c \geq 1$ is an absolute constant.

It is well known and easy to show that Theorem 6.8 is essentially tight, even for *nonadaptive* decision trees [**94**, Theorem 5.19]. The actual statement that we prove is more precise and takes into account the density parameter $\mathbf{P}[T(x) \neq 0]$; see Theorem 6.37 for details. With Theorem 6.8 in hand, all our main results (Theorem 6.1 and its corollaries) follow immediately by combining the new bound on the Fourier weight of decision trees with Tal's near-optimal bounds on the individual Fourier coefficients of $\mathcal{D}_{n,k,U}$.

Theorem 6.8 is of interest in its own right, independent of its use in this chapter to obtain optimal quantum-classical separations. The study of the Fourier spectrum has a variety of applications in theoretical computer science, including circuit complexity, learning theory, pseudorandom generators, and quantum computing. Even prior to Tal's work, the ℓ-Fourier weight of decision trees was studied for $\ell = 1$ by O'Donnell and Servedio [**95**], who proved the tight $O(\sqrt{d})$ bound and used it to give a polynomial-time learning algorithm for monotone decision trees. Fourier weight has been studied for various other classes of Boolean functions, including bounded-depth circuits, branching programs, low-degree polynomials over finite fields, and functions with bounded sensitivity; see the recent papers [**66, 133, 134, 46, 45**] and the references therein.

6.2. Generalized decision trees

Throughout this chapter, we assume decision trees to be perfect binary trees, with each internal node having two children and all leaves having the same depth. This convention is without loss of generality since a decision tree computing a given function f can be made into a perfect binary tree for f of the same depth, by querying dummy variables as necessary. We denote the variables of a decision tree by

$x_1, x_2, \ldots, x_n \in \{-1,1\}$, and identify the vertices of a decision tree in the natural manner with strings in $\{-1,1\}^*$. Thus, ε denotes the root of the tree, and a string $v \in \{-1,1\}^k$ denotes the vertex at depth k reached from the root by following the path $v_1 v_2 \ldots v_k$. Formally, a *decision tree* of depth d in Boolean variables $x_1, x_2, \ldots, x_n \in \{-1,1\}$ is a function T on $\{-1,1\}^{\leq d}$ with the following two properties.

(i) One has $T(v) \in \{1, 2, \ldots, n\}$ for every $v \in \{-1,1\}^{\leq d-1}$, with the interpretation that $T(v)$ is the index of the variable queried at the internal node found by following the path $v = v_1 v_2 v_3 \ldots$ from the root of the decision tree. We note that a variable cannot be queried twice on the same path, and therefore the d numbers $T(\varepsilon), T(v_1), T(v_1 v_2), \ldots, T(v_1 v_2 \ldots v_{d-1})$ are pairwise distinct for every $v \in \{-1,1\}^{d-1}$.

(ii) One has $T(v) \in \mathbb{R}[x_1, x_2, \ldots, x_n]$ for every $v \in \{-1,1\}^d$, with the interpretation that $T(v)$ is the label of the leaf reached by following the path $v = v_1 v_2 \ldots v_d$ from the root of the tree. Thus, every leaf is labeled with a real-valued polynomial in the input variables $x_1, x_2, \ldots, x_n$. At a given leaf $v \in \{-1,1\}^d$, the variables $x_{T(\varepsilon)}, x_{T(v_1)}, \ldots, x_{T(v_1 v_2 \ldots v_{d-1})}$ have been queried and therefore have fixed values. For this reason, we require $T(v)$ to be a real polynomial in variables other than $x_{T(\varepsilon)}, x_{T(v_1)}, \ldots, x_{T(v_1 v_2 \ldots v_{d-1})}$. We refer to a leaf $v \in \{-1,1\}^d$ as a *nonzero leaf* if $T(v)$ is not the zero polynomial. While we formally allow arbitrary real polynomials, the identity $x_i^2 = x_i$ effectively forces $T(v)$ for each $v \in \{-1,1\}^d$ to be multilinear.

Our formalism generalizes the traditional notion of a decision tree, where the leaf labels are restricted to the Boolean constants 0 and 1.

Proposition 6.9. *Let T be a given decision tree of depth d. Then the function $f\colon \{-1,1\}^n \to \mathbb{R}$ computed by T is given by*

$$f(x) = \sum_{v \in \{-1,1\}^d} T(v) \cdot \prod_{i=1}^{d} \frac{1 + v_i x_{T(v_1 v_2 \ldots v_{i-1})}}{2}. \tag{6.2.1}$$

We emphasize that $T(v)$ in this expression is a polynomial in $x_1, x_2, \ldots, x_n$ and not necessarily a constant value. In fact, the norm $\|\!|T(v)|\!\|$ for leaves v is a prominent quantity in this chapter.

Proof. For an input $x \in \{-1,1\}^n$ and a leaf $v \in \{-1,1\}^d$, the product

$$\prod_{i=1}^{d} \frac{1 + v_i x_{T(v_1 v_2 \ldots v_{i-1})}}{2}$$

evaluates to 1 if the input x reaches the leaf v in T, and evaluates to 0 otherwise. Recall that any given input x reaches precisely one leaf v, and the output of the tree on x is defined to be the corresponding polynomial $T(v) \in \mathbb{R}[x_1, x_2, \ldots, x_n]$ evaluated at x. Thus, (6.2.1) evaluates to $T(v)$ where v is the leaf reached by x. □

For a decision tree T of depth d, we let $\operatorname{dns}(T)$ denote the fraction of leaves in T with nonzero labels:

$$\operatorname{dns}(T) = \mathop{\mathbf{P}}_{v \in \{-1,1\}^d} [T(v) \neq 0].$$

We refer to this quantity as the *density* of T. Another important complexity measure is the *degree of* T, denoted $\deg(T)$ and defined as the maximum of the degrees of the polynomials $T(v) \in \mathbb{R}[x_1, x_2, \ldots, x_n]$ for $v \in \{-1,1\}^d$. Recall that the zero polynomial 0 is considered to have degree $-\infty$. For an internal node $v \in \{-1,1\}^{\leq d-1}$, we let T_v denote the subtree of T rooted at v. Thus, T_v is the tree of depth $d - |v|$ given by

$T_v(u) = T(vu)$ for all $u \in \{-1,1\}^{\leq d-|v|}$. The following fact is straightforward and well-known.

FACT 6.10. *Let T be a given decision tree of degree at most 0. Let $f\colon \{-1,1\}^n \to \mathbb{R}$ be the function computed by T. Then*

$$\mathop{\mathbf{P}}_{x\in\{-1,1\}^n}[f(x) \neq 0] = \operatorname{dns}(T).$$

Proof. Let d be the depth of T. Since T is a perfect binary tree, the fraction of inputs $x \in \{-1,1\}^n$ that reach any given leaf of T is exactly 2^{-d}. Therefore, the probability that a random input $x \in \{-1,1\}^n$ reaches a leaf with a nonzero label is precisely the fraction of leaves with nonzero labels, which is by definition $\operatorname{dns}(T)$. □

We will be working with special classes of trees described by several parameters. Specifically, we let $\mathcal{T}(n,d,p,k)$ denote the set of all trees in n Boolean variables $x_1, x_2, \ldots, x_n \in \{-1,1\}$ of depth d and density p such that for every leaf $v \in \{-1,1\}^d$, the label $T(v)$ is either the zero polynomial 0 or a homogeneous multilinear polynomial of degree k. We further define $\mathcal{T}^*(n,d,p,k)$ to be the set of all trees $T \in \mathcal{T}(n,d,p,k)$ that have the additional property that $T(v) \in \{0\} \cup \{\pm \prod_{i\in S} x_i \colon S \in \mathcal{P}_{n,k}\}$ for every leaf $v \in \{-1,1\}^d$. Thus, every nonzero leaf in a tree $T \in \mathcal{T}^*(n,d,p,k)$ is labeled with a signed monomial of degree k.

The Fourier spectrum of decision trees has been studied in several works, as discussed in the introduction. We will need the following special case of a result due to Tal [**135**, Theorem 7.5]. For completeness, we include our proof in Section 6.4.5.

THEOREM 6.11 (Tal). *Let $f\colon \{-1,1\}^n \to \{-1,0,1\}$ be given, $f \not\equiv 0$. Define $p = \mathbf{P}_{x\in\{-1,1\}^n}[f(x) \neq 0]$. Suppose that f can be computed by a depth-d decision tree.*

Then

$$\|\!|L_1 f\|\!| \le \binom{d}{1}^{1/2} Cp\sqrt{\ln\frac{e}{p}},$$
$$\|\!|L_2 f\|\!| \le \binom{d}{2}^{1/2} C^2p\sqrt{\ln\frac{e}{p}}\sqrt{\ln\frac{en}{p}},$$

where $C \ge 1$ is an absolute constant.

Tal states his result for functions $f\colon \{-1,1\}^n \to \{0,1\}$ rather than $f\colon \{-1,1\}^n \to \{-1,0,1\}$. But Theorem 6.11 follows immediately by writing $f = f^+ - f^-$, where $f^+, f^-\colon \{-1,1\}^n \to \{0,1\}$ are the positive and negative parts of f, and applying Tal's result separately to f^+ and f^-.

6.3. Elementary set families

As explained in the introduction, we obtain our Fourier weight bound by combining the Fourier coefficients of a decision tree into well-structured groups and bounding the sum of the absolute values in each group. In this section, we lay the combinatorial groundwork for this result by proving that $\mathcal{P}_{n,k}$ can be efficiently partitioned into what we call "elementary families." We start in Section 6.3.1 with some technical calculations. Section 6.3.2 formally defines elementary families and studies the associated complexity measure for representing general families as the disjoint union of elementary parts. Finally, Section 6.3.3 proves that our family of interest $\mathcal{P}_{n,k}$ has an efficient partition of this form.

6.3.1. A binomial recurrence. Our starting point is a technical calculation related to the entropy function.

Lemma 6.12. *There is an absolute constant $c \geq 1$ such that for all integers $k \geq 1$,*

$$\sum_{i=1}^{k-1} \left(\frac{k}{i}\right)^{i/2} \left(\frac{k}{k-i}\right)^{(k-i)/2} \frac{1}{\sqrt{i(k-i)}} \leq c\sqrt{\frac{2^k}{k}}.$$

Proof. To begin with,

$$\begin{aligned}
&\sum_{i=1}^{k-1} \left(\frac{k}{i}\right)^{i/2} \left(\frac{k}{k-i}\right)^{(k-i)/2} \frac{1}{\sqrt{i(k-i)}} \\
&\qquad = \sum_{i=1}^{k-1} \frac{2^{H(i/k)\cdot k/2}}{\sqrt{i(k-i)}} \\
&\qquad \leq 2^{k/2} \sum_{i=1}^{k-1} \exp\left(-k\left(\frac{i}{k}-\frac{1}{2}\right)^2\right) \cdot \frac{1}{\sqrt{i(k-i)}}, \qquad (6.3.1)
\end{aligned}$$

where the last step uses (2.1.2). Continuing,

$$\begin{aligned}
\sum_{i=1}^{\lceil k/4\rceil -1} \exp\left(-k\left(\frac{i}{k}-\frac{1}{2}\right)^2\right) \frac{1}{\sqrt{i(k-i)}} &\leq \sum_{i=1}^{\lceil k/4\rceil -1} \exp\left(-k\left(\frac{i}{k}-\frac{1}{2}\right)^2\right) \\
&\leq \sum_{i=1}^{\lceil k/4\rceil -1} e^{-k/16} \\
&< \frac{ke^{-k/16}}{4}. \qquad (6.3.2)
\end{aligned}$$

Symmetrically,

$$\sum_{i=\lfloor 3k/4\rfloor +1}^{k-1} \exp\left(-k\left(\frac{i}{k}-\frac{1}{2}\right)^2\right) \frac{1}{\sqrt{i(k-i)}} < \frac{ke^{-k/16}}{4}. \qquad (6.3.3)$$

Finally,

$$\begin{aligned}
\sum_{i=\lceil k/4\rceil}^{\lfloor 3k/4\rfloor} \exp\left(-k\left(\frac{i}{k}-\frac{1}{2}\right)^2\right)\frac{1}{\sqrt{i(k-i)}}&\\
&\leq \frac{4}{\sqrt{3}k}\sum_{i=\lceil k/4\rceil}^{\lfloor 3k/4\rfloor} \exp\left(-k\left(\frac{i}{k}-\frac{1}{2}\right)^2\right)\\
&\leq \frac{4}{\sqrt{3}k}\sum_{i=-\infty}^{\infty} \exp\left(-k\left(\frac{i}{k}-\frac{1}{2}\right)^2\right)\\
&\leq \frac{4}{\sqrt{3}k}+\frac{4}{\sqrt{3}k}\int_{-\infty}^{\infty} \exp\left(-k\left(\frac{x}{k}-\frac{1}{2}\right)^2\right)dx\\
&= \frac{4}{\sqrt{3}k}+\frac{4\sqrt{\pi}}{\sqrt{3k}}. \qquad (6.3.4)
\end{aligned}$$

Combining (6.3.1)–(6.3.4), we conclude that

$$\sum_{i=1}^{k-1}\left(\frac{k}{i}\right)^{i/2}\left(\frac{k}{k-i}\right)^{(k-i)/2}\frac{1}{\sqrt{i(k-i)}}\leq 2^{k/2}\left(\frac{ke^{-k/16}}{2}+\frac{4}{\sqrt{3}k}+\frac{4\sqrt{\pi}}{\sqrt{3k}}\right).$$

This settles the lemma for a large enough absolute constant $c \geq 1$. □

As an application of the previous lemma, we proceed to solve a key recurrence that we will need to study $\mathcal{P}_{n,k}$.

THEOREM 6.13. *Let* $N\colon \{1,2,4,8,16,\ldots\}\times\mathbb{Z}^+\to[0,\infty)$ *be any function that satisfies*

$$\begin{aligned}
N(n,k) &\leq \binom{n}{k}^{1/2} && \textit{if } \min\{n,k\}\leq 2,\\
N(n,k) &\leq 2N\left(\frac{n}{2},k\right)+\sum_{i=1}^{k-1}N\left(\frac{n}{2},i\right)N\left(\frac{n}{2},k-i\right) && \textit{if } \min\{n,k\}>2.
\end{aligned}$$

Let $c\geq 1$ *be the absolute constant from Lemma* 6.12. *Then for all* n,k,

$$N(n,k)\leq\frac{(2+\sqrt{2})^{k-1}c^{k-1}}{\sqrt{k}}\left(\frac{n}{k}\right)^{k/2}. \qquad (6.3.5)$$

Proof. The proof of (6.3.5) is by induction on the pair $(n,k)\in\{1,2,4,8,16,\ldots\}\times\mathbb{Z}^+$. For $\min\{n,k\}\le 2$, the claimed bound (6.3.5) is a weakening of $N(n,k)\le\binom{n}{k}^{1/2}$. This establishes the base case. For the inductive step, fix arbitrary $n\in\{4,8,16,32,\ldots\}$ and $k\ge 3$. Abbreviate $\alpha=2+\sqrt{2}$. Then

$$\begin{aligned}
N(n,k) &\le 2N\Big(\frac{n}{2},k\Big)+\sum_{i=1}^{k-1}N\Big(\frac{n}{2},i\Big)\,N\Big(\frac{n}{2},k-i\Big)\\
&\le 2\cdot\frac{(\alpha c)^{k-1}}{\sqrt{k}}\left(\frac{n}{2k}\right)^{k/2}\\
&\qquad+\sum_{i=1}^{k-1}\frac{(\alpha c)^{i-1}}{\sqrt{i}}\left(\frac{n}{2i}\right)^{i/2}\cdot\frac{(\alpha c)^{k-i-1}}{\sqrt{k-i}}\left(\frac{n}{2(k-i)}\right)^{(k-i)/2}\\
&=2\cdot\frac{(\alpha c)^{k-1}}{\sqrt{k}}\left(\frac{n}{2k}\right)^{k/2}\\
&\qquad+(\alpha c)^{k-2}\left(\frac{n}{2k}\right)^{k/2}\sum_{i=1}^{k-1}\frac{1}{\sqrt{i(k-i)}}\left(\frac{k}{i}\right)^{i/2}\left(\frac{k}{k-i}\right)^{(k-i)/2}\\
&\le 2\cdot\frac{(\alpha c)^{k-1}}{\sqrt{k}}\left(\frac{n}{2k}\right)^{k/2}+\frac{(\alpha c)^{k-2}c}{\sqrt{k}}\left(\frac{n}{k}\right)^{k/2}\\
&\le\frac{1}{\sqrt{2}}\cdot\frac{(\alpha c)^{k-1}}{\sqrt{k}}\left(\frac{n}{k}\right)^{k/2}+\frac{(\alpha c)^{k-2}c}{\sqrt{k}}\left(\frac{n}{k}\right)^{k/2}\\
&=\frac{(\alpha c)^{k-1}}{\sqrt{k}}\left(\frac{n}{k}\right)^{k/2},
\end{aligned}$$

where the second step applies the inductive hypothesis; the fourth step appeals to Lemma 6.12; and the fifth step uses $k\ge 3$. This completes the inductive step and thereby settles (6.3.5). □

6.3.2. The partition measure. For set families $\mathcal{A},\mathcal{B}\subseteq\mathcal{P}(\mathbb{Z})$, we define $\mathcal{A}*\mathcal{B}=\{A\cup B: A\in\mathcal{A}, B\in\mathcal{B}\}$. We collect basic properties of this operation in the proposition below.

PROPOSITION 6.14. *Let $\mathcal{A},\mathcal{B},\mathcal{C}\subseteq\mathcal{P}(\mathbb{Z})$ be given. Then:*

(i) $\mathcal{A} * \varnothing = \varnothing * \mathcal{A} = \varnothing$;

(ii) $\mathcal{A} * \{\varnothing\} = \{\varnothing\} * \mathcal{A} = \mathcal{A}$;

(iii) $(\mathcal{A} * \mathcal{B}) * \mathcal{C} = \mathcal{A} * (\mathcal{B} * \mathcal{C})$;

(iv) $\mathcal{A} * \mathcal{B} = \mathcal{B} * \mathcal{A}$;

(v) $(\mathcal{A} \cup \mathcal{B}) * \mathcal{C} = (\mathcal{A} * \mathcal{C}) \cup (\mathcal{B} * \mathcal{C})$.

Proof. All properties are immediate from the definition of the $*$ operation. □

We define an *integer interval* to be any finite set whose elements are consecutive integers, namely, $\{i, i+1, i+2, \ldots, j\}$ for some $i, j \in \mathbb{Z}$. As a special case, this includes the empty interval $\varnothing$. An *elementary family* is any family of the form

$$\mathcal{E} = \binom{I_1}{k_1} * \binom{I_2}{k_2} * \cdots * \binom{I_\ell}{k_\ell}, \tag{6.3.6}$$

where ℓ is a positive integer, $I_1, I_2, \ldots, I_\ell$ are pairwise disjoint integer intervals, and $k_1, k_2, \ldots, k_\ell \in \{0, 1, 2\}$. Trivial examples of elementary families are $\binom{\varnothing}{0} = \{\varnothing\}$ and $\binom{\varnothing}{1} = \varnothing$. Another example of an elementary family is the singleton family $\{A\}$ for any nonempty finite set $A \subseteq \mathbb{Z}$, using $\{A\} = \binom{\{a_1\}}{1} * \binom{\{a_2\}}{1} * \cdots * \binom{\{a_\ell\}}{1}$ where $a_1 < a_2 < \cdots < a_\ell$ are the distinct elements of A. We now define a partition measure that captures how efficiently a family can be partitioned into elementary families.

Definition 6.15 (Partition measure π). For any family $\mathcal{A} \subseteq \mathcal{P}(\{1, 2, \ldots, n\})$, define $\pi(\mathcal{A})$ to be the minimum

$$\sum_{i=1}^{N} |\mathcal{E}_i|^{1/2} \tag{6.3.7}$$

over all integers N and all elementary families $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_N$ that are pairwise disjoint and satisfy $\mathcal{E}_1 \cup \mathcal{E}_2 \cup \cdots \cup \mathcal{E}_N = \mathcal{A}$.

Straight from the definition,

$$\pi(\varnothing) = 0,$$

$$\pi(\{\varnothing\}) = 1.$$

More generally,

$$|\mathcal{A}|^{1/2} \leq \pi(\mathcal{A}) \leq |\mathcal{A}| \tag{6.3.8}$$

for every $\mathcal{A} \subseteq \mathcal{P}(\{1, 2, \ldots, n\})$. The upper bound here corresponds to the trivial partition $\mathcal{A} = \bigcup_{A \in \mathcal{A}} \{A\}$. The lower bound holds because (6.3.7) is no smaller than $(\sum |\mathcal{E}_i|)^{1/2} = |\mathcal{A}|^{1/2}$. The following four lemmas will be useful to us in analyzing the partition measure for families of interest.

Lemma 6.16. *Let $\mathcal{A}, \mathcal{B} \subseteq \mathcal{P}(\{1, 2, \ldots, n\})$ be given with $\mathcal{A} \cap \mathcal{B} = \varnothing$. Then*

$$\pi(\mathcal{A} \cup \mathcal{B}) \leq \pi(\mathcal{A}) + \pi(\mathcal{B}).$$

Proof. If $\mathcal{A} = \varnothing$ or $\mathcal{B} = \varnothing$, the claim is trivial. In the complementary case, let $\mathcal{A} = \mathcal{E}_1 \cup \cdots \cup \mathcal{E}_N$ and $\mathcal{B} = \mathcal{E}'_1 \cup \cdots \cup \mathcal{E}'_{N'}$ be partitions of $\mathcal{A}$ and $\mathcal{B}$, respectively, into elementary families. Then $\mathcal{A} \cup \mathcal{B} = (\mathcal{E}_1 \cup \cdots \cup \mathcal{E}_N) \cup (\mathcal{E}'_1 \cup \cdots \cup \mathcal{E}'_{N'})$ is a partition of $\mathcal{A} \cup \mathcal{B}$ into elementary families. □

Lemma 6.17. *Let $\mathcal{A} \subseteq \mathcal{P}(\{1, 2, \ldots, m\})$ and $\mathcal{B} \subseteq \mathcal{P}(\{m+1, m+2, \ldots, n\})$ be given, for some $1 \leq m < n$. Then*

$$\pi(\mathcal{A} * \mathcal{B}) \leq \pi(\mathcal{A})\, \pi(\mathcal{B}).$$

Proof. If $\mathcal{A} = \varnothing$ or $\mathcal{B} = \varnothing$, we have $\mathcal{A} * \mathcal{B} = \varnothing$ by Proposition 6.14 and therefore $\pi(\mathcal{A} * \mathcal{B}) = 0$. In the complementary case, let $\mathcal{A} = \mathcal{E}_1 \cup \cdots \cup \mathcal{E}_N$ and $\mathcal{B} = \mathcal{E}'_1 \cup \cdots \cup \mathcal{E}'_{N'}$ be partitions of $\mathcal{A}$ and $\mathcal{B}$, respectively, into elementary families for which $\pi(\mathcal{A})$ and

$\pi(\mathcal{B})$ are achieved. Then

$$\mathcal{A} * \mathcal{B} = \left(\bigcup_{i=1}^{N} \mathcal{E}_i \right) * \mathcal{B} = \bigcup_{i=1}^{N} (\mathcal{E}_i * \mathcal{B}) = \bigcup_{i=1}^{N} \bigcup_{j=1}^{N'} (\mathcal{E}_i * \mathcal{E}'_j), \tag{6.3.9}$$

where the last two steps use the distributivity and commutativity properties in Proposition 6.14. For any elementary families $\mathcal{E}_i \subseteq \mathcal{P}(\{1,2,\ldots,m\})$ and $\mathcal{E}'_j \subseteq \mathcal{P}(\{m+1, m+2, \ldots, n\})$, the family $\mathcal{E}_i * \mathcal{E}'_j \subseteq \mathcal{P}(\{1,2,\ldots,n\})$ is also elementary, with $|\mathcal{E}_i * \mathcal{E}'_j| = |\mathcal{E}_i|\,|\mathcal{E}'_j|$. Since all unions in (6.3.9) are disjoint, we obtain

$$\pi(\mathcal{A} * \mathcal{B}) \leq \sum_{i=1}^{N} \sum_{j=1}^{N'} |\mathcal{E}_i * \mathcal{E}'_j|^{1/2} = \sum_{i=1}^{N} \sum_{j=1}^{N'} |\mathcal{E}_i|^{1/2} |\mathcal{E}'_j|^{1/2} = \pi(\mathcal{A})\pi(\mathcal{B}). \qquad \square$$

For a set $A \subseteq \mathbb{Z}$ and an integer x, we define $A + x = \{a + x : a \in A\}$. Analogously, for a family $\mathcal{A} \subseteq \mathcal{P}(\mathbb{Z})$, we define $\mathcal{A} + x = \{A + x : A \in \mathcal{A}\}$. As one would expect, the partition measure is invariant under translation by an integer.

Lemma 6.18. *Let $\mathcal{A} \subseteq \mathcal{P}(\{1,2,\ldots,n\})$ be given. Then for all $x \in \mathbb{N}$,*

$$\pi(\mathcal{A}) = \pi(\mathcal{A} + x).$$

Proof. Consider an elementary family $\mathcal{E}$ of the form (6.3.6), where $I_1, I_2, \ldots, I_\ell$ are pairwise disjoint integer intervals and $k_1, k_2, \ldots, k_\ell \in \{0, 1, 2\}$. Then

$$\mathcal{E} + x = \binom{I_1 + x}{k_1} * \binom{I_2 + x}{k_2} * \cdots * \binom{I_\ell + x}{k_\ell}$$

is also an elementary family because the translated integer intervals $I_1 + x, I_2 + x, \ldots, I_\ell + x$ are pairwise disjoint. Thus, any partition $\mathcal{A} = \bigcup_{i=1}^{N} \mathcal{E}_i$ into elementary families gives an analogous partition $\mathcal{A} + x = \bigcup_{i=1}^{N} (\mathcal{E}_i + x)$ into elementary families, with $|\mathcal{E}_i + x| = |\mathcal{E}_i|$ for all i. $\square$

In general, $\mathcal{A} \subseteq \mathcal{B}$ does not imply $\pi(\mathcal{A}) \leq \pi(\mathcal{B})$. However, π enjoys the following monotonicity property.

LEMMA 6.19. *For any positive integers n, m, k with $n \leq m$,*

$$\pi(\mathcal{P}_{n,k}) \leq \pi(\mathcal{P}_{m,k}).$$

Proof. Consider an elementary family $\mathcal{E}$ of the form (6.3.6), where $I_1, I_2, \ldots, I_\ell$ are pairwise disjoint integer intervals and $k_1, k_2, \ldots, k_\ell \in \{0, 1, 2\}$. Then

$$\mathcal{E} \cap \mathcal{P}(\{1, 2, \ldots, n\}) = \binom{I_1 \cap \{1, 2, \ldots, n\}}{k_1} * \cdots * \binom{I_\ell \cap \{1, 2, \ldots, n\}}{k_\ell}$$

is also an elementary family because the integer intervals $I_j \cap \{1, 2, \ldots, n\}$ for $j = 1, 2, \ldots, \ell$ are pairwise disjoint. Thus, any partition $\mathcal{P}_{m,k} = \bigcup_{i=1}^{N} \mathcal{E}_i$ into elementary families gives an analogous partition for $\mathcal{P}_{n,k}$:

$$\begin{aligned} \mathcal{P}_{n,k} &= \mathcal{P}_{m,k} \cap \mathcal{P}(\{1, 2, \ldots, n\}) \\ &= \bigcup_{i=1}^{N} \mathcal{E}_i \cap \mathcal{P}(\{1, 2, \ldots, n\}). \end{aligned}$$

Moreover, the elementary families in the new partition obey $|\mathcal{E}_i \cap \mathcal{P}(\{1, 2, \ldots, n\})| \leq |\mathcal{E}_i|$ for all i. □

6.3.3. An efficient partition for $\mathcal{P}_{n,k}$. Our analysis of the Fourier spectrum of decision trees relies on the partition measure of the family $\mathcal{P}_{n,k}$. Recall from (6.3.8) that

$$\pi(\mathcal{P}_{n,k}) \geq \binom{n}{k}^{1/2}.$$

We will now prove that this lower bound is tight up to a factor of $2^{O(k)}$, by combining Lemmas 6.16–6.19 with the recurrence solved in Theorem 6.13.

THEOREM 6.20. *Let $c \geq 1$ be the absolute constant from Lemma* 6.12. *Then for all positive integers n and k,*

$$\pi(\mathcal{P}_{n,k}) \leq \frac{(2+\sqrt{2})^{k-1}c^{k-1}}{\sqrt{k}} \left(\frac{2n}{k}\right)^{k/2}. \tag{6.3.10}$$

Proof. We first treat the case when n is a power of 2. If $k \leq 2$, the family $\mathcal{P}_{n,k}$ is elementary to start with. As a result,

$$\pi(\mathcal{P}_{n,k}) \leq \binom{n}{k}^{1/2}, \qquad\qquad k \leq 2. \tag{6.3.11}$$

If $n \leq 2$, the family $\mathcal{P}_{n,k}$ is empty unless $k \leq 2$. Therefore, again

$$\pi(\mathcal{P}_{n,k}) \leq \binom{n}{k}^{1/2}, \qquad\qquad n \leq 2. \tag{6.3.12}$$

For $n, k \geq 3$, we have

$$\begin{aligned}
\pi(\mathcal{P}_{n,k}) &= \pi\left(\bigcup_{i=0}^{k}\left(\binom{\{1,2,\ldots,n/2\}}{i} * \binom{\{n/2+1, n/2+2, \ldots, n\}}{k-i}\right)\right) \\
&\leq \sum_{i=0}^{k} \pi\left(\binom{\{1,2,\ldots,n/2\}}{i} * \binom{\{n/2+1, n/2+2, \ldots, n\}}{k-i}\right) \\
&\leq \sum_{i=0}^{k} \pi\left(\binom{\{1,2,\ldots,n/2\}}{i}\right) \pi\left(\binom{\{n/2+1, n/2+2, \ldots, n\}}{k-i}\right) \\
&= \sum_{i=0}^{k} \pi(\mathcal{P}_{n/2,i})\, \pi\left(\mathcal{P}_{n/2,k-i} + \frac{n}{2}\right) \\
&= \sum_{i=0}^{k} \pi(\mathcal{P}_{n/2,i})\, \pi(\mathcal{P}_{n/2,k-i}) \\
&= 2\pi(\mathcal{P}_{n/2,k}) + \sum_{i=1}^{k-1} \pi(\mathcal{P}_{n/2,i})\, \pi(\mathcal{P}_{n/2,k-i}),
\end{aligned} \tag{6.3.13}$$

where the second, third, and fifth steps apply Lemmas 6.16, 6.17, and 6.18, respectively, and the last step uses $\pi(\{\varnothing\}) = 1$.

The recurrence relations (6.3.11)–(6.3.13) show that the hypothesis of Theorem 6.13 is satisfied for the function $N(n,k) := \pi(\mathcal{P}_{n,k})$. As a result, Theorem 6.13 implies that

$$\pi(\mathcal{P}_{n,k}) \leq \frac{(2+\sqrt{2})^{k-1}c^{k-1}}{\sqrt{k}} \left(\frac{n}{k}\right)^{k/2}$$

for any $n \in \{1,2,4,8,16,\ldots\}$ and $k \geq 1$. This upper bound in turn implies (6.3.10) for any $n \geq 1$ and $k \geq 1$:

$$\begin{aligned}\pi(\mathcal{P}_{n,k}) &\leq \pi(\mathcal{P}_{2^{\lceil \log n\rceil},k})\\ &\leq \frac{(2+\sqrt{2})^{k-1}c^{k-1}}{\sqrt{k}} \left(\frac{2^{\lceil \log n\rceil}}{k}\right)^{k/2}\\ &\leq \frac{(2+\sqrt{2})^{k-1}c^{k-1}}{\sqrt{k}} \left(\frac{2n}{k}\right)^{k/2},\end{aligned}$$

where the first step uses Lemma 6.19. □

6.4. Fourier spectrum of decision trees

This section is devoted to the proof of our main result on the Fourier spectrum of decision trees. Stated in its simplest terms, our result shows that for any function $f\colon \{-1,1\}^n \to \{-1,0,1\}$ computable by a decision tree of depth d, the sum of the absolute values of the Fourier coefficients of order k is at most

$$C^k\sqrt{\binom{d}{k}(1+\ln n)^{k-1}},$$

where $C \geq 1$ is an absolute constant that does not depend on n,d,k. Sections 6.4.1–6.4.3 focus on partitioning the Fourier spectrum of f into highly structured parts and analyzing each in isolation. Sections 6.4.4 and 6.4.6 then recombine these pieces using the machinery of elementary families.

6.4.1. Slicing the tree. Let T be a given decision tree of depth d in Boolean variables $x_1, x_2, \ldots, x_n$. For a set family $\mathcal{S} \subseteq \mathcal{P}(\{1, 2, \ldots, d\})$, we define a real function $T|_{\mathcal{S}} \colon \{-1,1\}^n \to \mathbb{R}$ by

$$T|_{\mathcal{S}}(x) = \sum_{S \in \mathcal{S}} \sum_{v \in \{-1,1\}^d} T(v) \cdot 2^{-d} \prod_{i \in S} v_i x_{T(v_1 v_2 \ldots v_{i-1})}. \tag{6.4.1}$$

A straightforward but crucial observation is that $T|_{\mathcal{S}}$ is additive with respect to $\mathcal{S}$, in the following sense.

Proposition 6.21. *Let T be a depth-d decision tree. Let $\mathcal{S}', \mathcal{S}'' \subseteq \mathcal{P}(\{1, 2, \ldots, d\})$ be set families with $\mathcal{S}' \cap \mathcal{S}'' = \varnothing$. Then*

$$T|_{\mathcal{S}' \cup \mathcal{S}''} = T|_{\mathcal{S}'} + T|_{\mathcal{S}''}.$$

Proof. Immediate by taking $\mathcal{S} = \mathcal{S}' \cup \mathcal{S}''$ in the defining equation (6.4.1). □

The relevance of (6.4.1) to the Fourier spectrum of decision trees is borne out by the following lemma.

Lemma 6.22. *Let T be a decision tree of depth d and degree at most* 0, *computing a function $f \colon \{-1,1\}^n \to \mathbb{R}$. Then*

$$L_k f = T|_{\mathcal{P}_{d,k}}, \qquad\qquad k = 0, 1, 2, \ldots, n.$$

Proof. By Proposition 6.9,

$$\begin{aligned}
f(x) &= \sum_{v\in\{-1,1\}^d} T(v)\cdot\prod_{i=1}^{d}\frac{1+v_i x_{T(v_1v_2\ldots v_{i-1})}}{2}\\
&= \sum_{v\in\{-1,1\}^d} T(v)\cdot 2^{-d}\sum_{S\subseteq\{1,2,\ldots,d\}}\prod_{i\in S} v_i x_{T(v_1v_2\ldots v_{i-1})}\\
&= \sum_{k=0}^{d}\sum_{S\in\mathcal{P}_{d,k}}\sum_{v\in\{-1,1\}^d} T(v)\cdot 2^{-d}\prod_{i\in S} v_i x_{T(v_1v_2\ldots v_{i-1})}. \qquad (6.4.2)
\end{aligned}$$

Since $\deg(T)\le 0$, the coefficients $T(v)$ for $v\in\{-1,1\}^d$ are real numbers. Moreover, for any $v\in\{-1,1\}^d$ and $S\subseteq\{1,2,\ldots,d\}$, the definition of a decision tree ensures that the product $\prod_{i\in S} v_i x_{T(v_1v_2\ldots v_{i-1})}$ is a signed monomial of degree $|S|$. We conclude from (6.4.2) that the degree-k homogeneous part of f is

$$\begin{aligned}
L_k f &= \sum_{S\in\mathcal{P}_{d,k}}\sum_{v\in\{-1,1\}^d} T(v)\cdot 2^{-d}\prod_{i\in S} v_i x_{T(v_1v_2\ldots v_{i-1})}\\
&= T|_{\mathcal{P}_{d,k}}.
\end{aligned}$$

In particular, $L_k f = 0$ for $k\ge d+1$. □

Looking ahead, much of our analysis of the Fourier spectrum of decision trees T focuses on $T|_{\mathcal{E}}$ for elementary families $\mathcal{E}\subseteq\mathcal{P}_{d,k}$. This analysis proceeds by induction, with the following lemma required as part of the inductive step.

Lemma 6.23. *Let $T\in\mathcal{T}(n,d,p,k)$ be a given decision tree and $\mathcal{S}\subseteq\mathcal{P}(\{1,2,\ldots,d\})$. Define $m=\max_{v\in\{-1,1\}^d}\|T(v)\|$. Then for each $i=1,2,\ldots,\binom{n}{k}$, there is a real $0\le$*

$p_i \leq 1$ *and a decision tree* $U_i \in \mathcal{T}^*(n, d, p_i, 0)$ *such that*

$$p = \sum_{i=1}^{\binom{n}{k}} p_i,$$

$$\|T|_{\mathcal{S}}\| \leq m \sum_{i=1}^{\binom{n}{k}} \|U_i|_{\mathcal{S}}\|.$$

Proof. Let $\phi = \sum_{S \subseteq \{1,2,\ldots,n\}} \hat{\phi}(S)\chi_S$ be an arbitrary nonzero polynomial with $\|\phi\| \leq 1$. Consider the random variable $X \in \{\pm\chi_S : \hat{\phi}(S) \neq 0\}$ distributed according to

$$\mathbf{P}[X = \sigma\chi_S] = \frac{|\hat{\phi}(S)|}{\|\phi\|} \left(\frac{1}{2} + \frac{\|\phi\|}{2} \cdot \sigma \operatorname{sgn} \hat{\phi}(S) \right)$$

for all $\sigma \in \{-1, 1\}$ and $S \subseteq \{1, 2, \ldots, n\}$. Then

$$\begin{aligned} \mathbf{E}\, X &= \sum_{S \subseteq \{1,2,\ldots,n\}} \sum_{\sigma \in \{-1,1\}} \sigma\chi_S \cdot \frac{|\hat{\phi}(S)|}{\|\phi\|} \left(\frac{1}{2} + \frac{\|\phi\|}{2} \cdot \sigma \operatorname{sgn} \hat{\phi}(S) \right) \\ &= \sum_{S \subseteq \{1,2,\ldots,n\}} \chi_S \cdot \frac{|\hat{\phi}(S)|}{\|\phi\|} \cdot \|\phi\| \cdot \operatorname{sgn} \hat{\phi}(S) \\ &= \phi(x). \end{aligned}$$

In conclusion, ϕ can be viewed as the *expected value* of a random variable $X \in \{\pm\chi_S : \hat{\phi}(S) \neq 0\}$.

We may assume that T has at least one nonzero leaf, since otherwise the lemma holds trivially with $p_1 = p_2 = \cdots = p_{\binom{n}{k}} = p = 0$. The previous paragraph implies that for every leaf $v \in \{-1, 1\}^d$ with $T(v) \neq 0$, the polynomial $T(v)/m$ is the expected value of a random variable X_v whose support is contained in the set of the nonzero degree-k monomials of $T(v)$ with ±1 coefficients. The joint distribution of the X_v is immaterial for our purposes, but for concreteness let us declare them to be independent. Then

$$\begin{aligned} T|_{\mathcal{S}}(x) &= m\sum_{S\in\mathcal{S}}\sum_{v\in\{-1,1\}^d} \frac{T(v)}{m}\cdot 2^{-d}\prod_{i\in S} v_i x_{T(v_1v_2\dots v_{i-1})} \\ &= m\sum_{S\in\mathcal{S}}\sum_{\substack{v\in\{-1,1\}^d:\\ T(v)\neq 0}} \mathbf{E}[X_v]\cdot 2^{-d}\prod_{i\in S} v_i x_{T(v_1v_2\dots v_{i-1})} \\ &= m\,\mathbf{E}\left[\sum_{S\in\mathcal{S}}\sum_{\substack{v\in\{-1,1\}^d:\\ T(v)\neq 0}} X_v\cdot 2^{-d}\prod_{i\in S} v_i x_{T(v_1v_2\dots v_{i-1})}\right]. \end{aligned}$$

Applying Proposition 2.2,

$$\| T|_{\mathcal{S}} \| \leq m\,\mathbf{E}\left\|\sum_{S\in\mathcal{S}}\sum_{\substack{v\in\{-1,1\}^d:\\ T(v)\neq 0}} X_v\cdot 2^{-d}\prod_{i\in S} v_i x_{T(v_1v_2\dots v_{i-1})}\right\|. \tag{6.4.3}$$

In the last expression, each random variable X_v is a signed monomial of degree k that does not contain any of the variables $x_{T(\varepsilon)}, x_{T(v_1)}, \dots, x_{T(v_1v_2\dots v_{d-1})}$ queried along the path from the root to v. Therefore, the expectation in (6.4.3) is over $\| U|_{\mathcal{S}}\|$ for some trees $U \in \mathcal{T}^*(n,d,p,k)$. We conclude that there is a fixed decision tree $U \in \mathcal{T}^*(n,d,p,k)$ with

$$\| T|_{\mathcal{S}} \| \leq m\, \| U|_{\mathcal{S}} \|. \tag{6.4.4}$$

Finally, decompose

$$U|_{\mathcal{S}} = \sum_{S\in\mathcal{P}_{n,k}} U_S|_{\mathcal{S}}\cdot \chi_S,$$

where U_S is the depth-d decision tree given by

$$U_S(v) = \begin{cases} U(v) & \text{if } |v| \leq d-1, \\ -1 & \text{if } |v| = d \text{ and } U(v) = -\chi_S, \\ 1 & \text{if } |v| = d \text{ and } U(v) = \chi_S, \\ 0 & \text{otherwise.} \end{cases}$$

In other words, U_S is the decision tree obtained from U by setting to 1 every leaf labeled χ_S, setting to -1 every leaf labeled $-\chi_S$, and setting all other leaves to 0. It is clear that the densities of the U_S sum to the density of U. We conclude that $U_S \in \mathcal{T}^*(n, d, p_S, 0)$ for some reals $0 \leq p_S \leq 1$ with $\sum_{S \in \mathcal{P}_{n,k}} p_S = p$. Moreover,

$$\begin{aligned} \|T|_{\mathcal{S}}\| &\leq m \, \|U|_{\mathcal{S}}\| \\ &\leq m \sum_{S \in \mathcal{P}_{n,k}} \|U_S|_{\mathcal{S}} \cdot \chi_S\| \\ &\leq m \sum_{S \in \mathcal{P}_{n,k}} \|U_S|_{\mathcal{S}}\|, \end{aligned}$$

where the first step is a restatement of (6.4.4); the second step applies Proposition 2.2; and the last step is justified by Proposition 2.3. In summary, the decision trees $U_1, U_2, \ldots, U_{\binom{n}{k}}$ in the statement of the lemma can be taken to be the U_S, in arbitrary order. $\square$

6.4.2. Analytic preliminaries.

For positive integers m and k, define

$$\Lambda_{m,k}(p) = \begin{cases} 0 & \text{if } p = 0, \\ p\sqrt{\left(\dfrac{1}{k}\ln\dfrac{e^k m^{k-1}}{p}\right)^k} & \text{if } 0 < p \leq 1/m, \\ p\sqrt{\left(\ln\dfrac{e}{p}\right)(\ln em)^{k-1}} & \text{if } 1/m < p \leq 1. \end{cases}$$

Our bound for the Fourier spectrum of decision trees is in terms of this function. As preparation for our main result, we now collect the analytic properties of $\Lambda_{m,k}$ that we will need.

LEMMA 6.24. *Let m and k be any positive integers. Then:*

(i) *$\Lambda_{m,k}$ is continuous on $[0,1]$;*

(ii) *$\Lambda_{m,k}$ is monotonically increasing on $[0,1]$;*

(iii) *$\Lambda_{m,k}$ is concave on $[0,1]$.*

Proof. (i) The continuity on $(0,1/m)\cup(1/m,1]$ is immediate. The continuity at $p=0$ and $p=1/m$ follows by examining the one-sided limits at those points, which are 0 and $(\ln em)^{k/2}/m$, respectively.

(ii) Considering the derivative $\Lambda'_{m,k}$ separately on $(0,1/m)$ and $(1/m,1]$, one finds in both cases that the derivative is positive:

$$\Lambda'_{m,k}(p)=\begin{cases}\sqrt{\left(\dfrac{1}{k}\ln\dfrac{e^k m^{k-1}}{p}\right)^k}\left(1-\dfrac{k}{2\ln(e^k m^{k-1}/p)}\right) & \text{if } 0<p<1/m,\\[2ex] \left(\sqrt{\ln\dfrac{e}{p}}-\dfrac{1}{2\sqrt{\ln(e/p)}}\right)\sqrt{(\ln em)^{k-1}} & \text{if } 1/m<p\leq 1.\end{cases}$$

Since $\Lambda_{m,k}$ is continuous on $[0,1]$, it follows that $\Lambda_{m,k}$ is monotonically increasing on $[0,1]$.

(iii) The one-sided derivatives of $\Lambda_{m,k}$ at $p=1/m$ are both $(\ln em)^{\frac{k-2}{2}}\ln(\sqrt{e}m)$. Along with the calculations in (ii), this shows that $\Lambda_{m,k}$ is continuously differentiable on $(0,1]$. The formulas in (ii) further reveal that $\Lambda'_{m,k}$ is monotonically decreasing on $(0,1/m)$ and on $(1/m,1]$. By the continuity of $\Lambda'_{m,k}$ on $(0,1]$, we conclude that $\Lambda'_{m,k}$ is monotonically decreasing on $(0,1]$, which in turn makes $\Lambda_{m,k}$ concave on $(0,1]$. Since $\Lambda_{m,k}$ is continuous at 0, we conclude that $\Lambda_{m,k}$ is concave on the entire interval $[0,1]$. □

The function $\Lambda_{m,k}$ arises as the solution to a natural optimization problem, which we now describe.

LEMMA 6.25. *Let m and k be positive integers. Then for $0 < p \leq 1$,*

$$\Lambda_{m,k}(p) = p \max \left\{ \prod_{i=1}^{k} \sqrt{\ln ex_i} : x_i \geq 1 \textit{ and } x_1 x_2 \dots x_i \leq \frac{m^{i-1}}{p} \textit{ for all } i \right\}. \tag{6.4.5}$$

Proof. For $k = 1$, the left-hand side and right-hand side are clearly $p\sqrt{\ln(e/p)}$. In what follows, we treat the complementary case $k \geq 2$.

For $0 < p \leq 1/m$, the upper bound in (6.4.5) follows by taking $x_1 = x_2 = \cdots = x_k = (m^{k-1}/p)^{1/k}$. For $1/m < p \leq 1$, the upper bound follows by setting $x_1 = 1/p$ and $x_2 = \cdots = x_k = m$.

For the lower bound in (6.4.5), fix reals $x_1, x_2, \dots, x_k \geq 1$ with $x_1 \leq 1/p$ and $x_1 x_2 \dots x_k \leq m^{k-1}/p$. Then

$$\begin{aligned} \sqrt{\ln ex_1} \cdot \prod_{i=2}^{k} \sqrt{\ln ex_i} &\leq \sqrt{\ln ex_1} \left(\frac{1}{k-1} \ln e^{k-1} x_2 \dots x_k \right)^{(k-1)/2} \\ &\leq \sqrt{\ln ex_1} \left(\frac{1}{k-1} \ln \frac{e^{k-1} m^{k-1}}{p x_1} \right)^{(k-1)/2}, \end{aligned} \tag{6.4.6}$$

where the first step applies the AM–GM inequality. Elementary calculus shows that (6.4.6) as a function of x_1 is monotonically increasing on $[1, (m^{k-1}/p)^{1/k}]$ and monotonically decreasing on $[(m^{k-1}/p)^{1/k}, m^{k-1}/p]$. Recalling that $1 \leq x_1 \leq 1/p$, we conclude that (6.4.6) is maximized at

$$x_1 = \min\left(\left(\frac{m^{k-1}}{p}\right)^{1/k}, \frac{1}{p}\right)$$
$$= \begin{cases} (m^{k-1}/p)^{1/k} & \text{if } 0 < p \leq 1/m, \\ 1/p & \text{if } 1/m < p \leq 1. \end{cases}$$

Making this substitution shows that (6.4.6) does not exceed $\Lambda_{m,k}(p)$. □

This optimization view of $\Lambda_{m,k}$ implies a host of useful facts that would be bothersome to prove directly. We state them as corollaries below.

Corollary 6.26. *Let m and k be positive integers. Then for all $p, q \in [0,1]$,*

$$q\Lambda_{m,k}(p) \leq \Lambda_{m,k}(pq).$$

Proof. If $p = 0$ or $q = 0$, the left-hand side and right-hand side both vanish. If $p, q \in (0,1]$, the claim can be equivalently stated as $\Lambda_{m,k}(p)/p \leq \Lambda_{m,k}(pq)/pq$, which in turn amounts to saying that $\Lambda_{m,k}(p)/p$ is monotonically nonincreasing in $p \in (0,1]$. This monotonicity is immediate from Lemma 6.25. □

Corollary 6.27. *Let m, k, ℓ be positive integers. Then for all $p, q \in [0,1]$,*

$$\Lambda_{m,k}(p)\,\Lambda_{m,\ell}\Big(\frac{q}{m}\Big) \leq \frac{\Lambda_{m,k+\ell}(pq)}{m}.$$

Proof. If $p = 0$ or $q = 0$, the left-hand side and right-hand side both vanish. In what follows, we treat $p, q \in (0,1]$. By Lemma 6.25,

$$\Lambda_{m,k}(p)\,\Lambda_{m,\ell}\Big(\frac{q}{m}\Big) = \frac{pq}{m}\max\left\{\prod_{i=1}^{k+\ell}\sqrt{\ln ex_i}\right\}, \tag{6.4.7}$$

where the maximum is over all $x_1, x_2, \dots, x_{k+\ell} \geq 1$ such that

$$x_1 x_2 \dots x_i \leq \frac{m^{i-1}}{p}, \qquad i = 1, 2, \dots, k, \tag{6.4.8}$$

$$x_{k+1} x_{k+2} \dots x_i \leq \frac{m^{i-k-1}}{q/m}, \qquad i = k+1, \dots, k+\ell. \tag{6.4.9}$$

Equations (6.4.8) and (6.4.9) imply that the maximum in (6.4.7) is over $x_1, x_2, \dots, x_{k+\ell} \geq 1$ that satisfy, among other things, $x_1 x_2 \dots x_i \leq m^{i-1}/(pq)$ for $i = 1, 2, \dots, k+\ell$. Now Lemma 6.25 implies that the right-hand side of (6.4.7) is at most $\Lambda_{m,k+\ell}(pq)/m$. □

Corollary 6.28. *Let m and k be positive integers. Then for all $p \in [0,1]$,*

$$\Lambda_{m,k}(p) \leq \sqrt{2^k p} \cdot \Lambda_{m,k}(\sqrt{p}). \tag{6.4.10}$$

Proof. For $p = 0$, the left-hand side and right-hand side both vanish. For $p \in (0,1]$, we have:

$$\begin{aligned}
\Lambda_{m,k}(p) &= p \max \left\{ \prod_{i=1}^{k} \sqrt{\ln e x_i} : x_i \geq 1 \text{ and } x_1 x_2 \dots x_i \leq \frac{m^{i-1}}{p} \text{ for all } i \right\} \\
&\leq p \max \left\{ \prod_{i=1}^{k} \sqrt{\ln e x_i^2} : x_i \geq 1 \text{ and } x_1 x_2 \dots x_i \leq \frac{m^{i-1}}{\sqrt{p}} \text{ for all } i \right\} \\
&\leq \sqrt{2^k}\, p \max \left\{ \prod_{i=1}^{k} \sqrt{\ln e x_i} : x_i \geq 1 \text{ and } x_1 x_2 \dots x_i \leq \frac{m^{i-1}}{\sqrt{p}} \text{ for all } i \right\} \\
&= \sqrt{2^k p} \cdot \Lambda_{m,k}(\sqrt{p}),
\end{aligned}$$

where the first and last steps use Lemma 6.25. □

6.4.3. Contiguous intervals. We have reached a focal point of this chapter, where we analyze $T|_{\mathcal{E}}$ for arbitrary decision trees T and "canonical" elementary families $\mathcal{E}$. The families that we allow are those of the form

$$\mathcal{E} = \binom{I_1}{k_1} * \binom{I_2}{k_2} * \cdots * \binom{I_\ell}{k_\ell},$$

where $k_1, k_2, \ldots, k_\ell \in \{1, 2\}$ and the integer intervals $I_1, I_2, \ldots, I_\ell$ form a partition of $\{1, 2, \ldots, d\}$ with d being the depth of T. The proof proceeds by induction on ℓ, with Lemmas 6.22, 6.23, and the analytic properties of $\Lambda_{m,k}$ applied in the inductive step. We will later generalize this result to arbitrary elementary families $\mathcal{E}$ and, from there, to all of $\mathcal{P}_{d,k}$ via the results of Section 6.3.

THEOREM 6.29. *Let $T \in \mathcal{T}^*(n, d, p, 0)$ be given, for some $0 \le p \le 1$ and integers $n, d \ge 1$. Let $\ell \ge 1$. Let $I_1, I_2, \ldots, I_\ell$ be pairwise disjoint integer intervals with $I_1 \cup I_2 \cup \cdots \cup I_\ell = \{1, 2, \ldots, d\}$, and let $k_1, k_2, \ldots, k_\ell \in \{1, 2\}$. Abbreviate $k = k_1 + k_2 + \cdots + k_\ell$. Then*

$$\left\| T|_{\binom{I_1}{k_1} * \binom{I_2}{k_2} * \cdots * \binom{I_\ell}{k_\ell}} \right\| \le 2C^k\, 12^{\ell-1} \Lambda_{n^2,k}(p) \prod_{i=1}^{\ell} \binom{|I_i|}{k_i}^{1/2}, \qquad (6.4.11)$$

where $C \ge 1$ is the absolute constant from Theorem 6.11.

Proof. The proof is by induction on ℓ. The base case $\ell = 1$ corresponds to $I_1 = \{1, 2, \ldots, d\}$. Let $f \colon \{-1,1\}^n \to \{-1, 0, 1\}$ be the function computed by T. If $f \equiv 0$, we have $T|_{\binom{I_1}{k_1}} \equiv 0$ and the bound holds trivially. In the complementary case $f \not\equiv 0$, recall from Fact 6.10 that

$$\mathop{\mathbf{P}}_{x \in \{-1,1\}^n} [f(x) \ne 0] = p. \qquad (6.4.12)$$

Then

$$\begin{aligned}
\|T|_{\binom{I_1}{k_1}}\| &= \|L_{k_1} f\| \\
&\leq \binom{|I_1|}{k_1}^{1/2} C^{k_1} p \prod_{i=1}^{k_1} \sqrt{\ln \frac{en^{i-1}}{p}} \\
&\leq \binom{|I_1|}{k_1}^{1/2} \cdot 2C^{k_1} p \prod_{i=1}^{k_1} \sqrt{\ln \frac{en^{i-1}}{\sqrt{p}}} \\
&\leq \binom{|I_1|}{k_1}^{1/2} \cdot 2C^{k_1} \Lambda_{n^2,k_1}(p) \\
&= \binom{|I_1|}{k_1}^{1/2} \cdot 2C^{k} \Lambda_{n^2,k}(p),
\end{aligned}$$

where the first step is valid by Lemma 6.22; the second step uses Theorem 6.11 along with (6.4.12) and $k_1 \leq 2$; and the fourth step applies Lemma 6.25. This settles the base case.

We now turn to the inductive step, $\ell \geq 2$. If $k_j > |I_j|$ for some j, then

$$T|_{\binom{I_1}{k_1}*\binom{I_2}{k_2}*\cdots*\binom{I_\ell}{k_\ell}} = T|_\varnothing = 0,$$

and the claimed bound holds trivially. We may therefore assume that $k_j \leq |I_j|$ for every $j = 1, 2, \ldots, \ell$. This means in particular that the intervals $I_1, I_2, \ldots, I_\ell$ are nonempty. Furthermore, by renumbering the intervals if necessary, we may assume that $I_1 < I_2 < \cdots < I_\ell$. Put $d' = \max I_{\ell-1}$, so that $I_\ell = \{d'+1, d'+2, \ldots, d\}$. Abbreviate

$$\begin{aligned}
\mathcal{S}' &= \binom{I_1}{k_1} * \binom{I_2}{k_2} * \cdots * \binom{I_{\ell-1}}{k_{\ell-1}}, \\
\mathcal{S} &= \mathcal{S}' * \binom{I_\ell}{k_\ell}.
\end{aligned}$$

For $j = 0, 1, 2, \ldots$, define a depth-d' decision tree T'_j by

$$T'_j(v) = \begin{cases} T(v) & \text{if } v \in \{-1,1\}^{\leq d'-1}, \\ T_v|_{\binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}} & \text{if } v \in \{-1,1\}^{d'} \text{ and } \operatorname{dns}(T_v) \in (3^{-j-1}, 3^{-j}]\,, \\ 0 & \text{otherwise.} \end{cases}$$

Observe that T'_j is a valid decision tree in that for every leaf $v \in \{-1,1\}^{d'}$, the label $T'_j(v) \in \mathbb{R}[x_1, x_2, \ldots, x_n]$ is a function that does not depend on any of the variables

$$x_{T(\varepsilon)}, x_{T(v_1)}, x_{T(v_1 v_2)}, \ldots, x_{T(v_1 v_2 \ldots v_{d'-1})} \tag{6.4.13}$$

queried along the path from the root to v. Indeed, recall from Lemma 6.22 that $T_v|_{\binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}}$ is the k_ℓ-th homogeneous part of the function computed by the subtree T_v, which by definition does not use any of the variables (6.4.13). We also note that all but finitely many of the trees $T_0, T_1, T_2, \ldots$ are identically zero; however, working with the infinite sequence is more convenient from the point of view of notation and calculations.

The weighted densities of $T'_0, T'_1, T'_2, \ldots$ are given by

$$\begin{aligned} \sum_{j=0}^{\infty} 3^{-j} \operatorname{dns}(T'_j) &= \sum_{j=0}^{\infty} 3^{-j} \mathop{\mathbf{P}}_{v \in \{-1,1\}^{d'}} [T'_j(v) \neq 0] \\ &\leq \sum_{j=0}^{\infty} 3^{-j} \mathop{\mathbf{P}}_{v \in \{-1,1\}^{d'}} [3^{-j-1} < \operatorname{dns}(T_v) \leq 3^{-j}] \\ &\leq 3 \mathop{\mathbf{E}}_{v \in \{-1,1\}^{d'}} \operatorname{dns}(T_v) \\ &= 3 \operatorname{dns}(T) \\ &= 3p. \end{aligned} \tag{6.4.14}$$

The relevance of T'_j to our analysis of $T|_{\mathcal{S}}$ is clear from the following claims, whose proofs we will present shortly.

Claim 6.30. *$T|_{\mathcal{S}} = \sum_{j=0}^{\infty} T'_j|_{\mathcal{S}'}$.*

Claim 6.31. *For $j = 0, 1, 2, \ldots,$ one has*

$$\|T'_j|_{\mathcal{S}'}\| \leq 8C^k\, 12^{\ell-2} \binom{|I_1|}{k_1}^{1/2} \cdots \binom{|I_\ell|}{k_\ell}^{1/2} \cdot \sqrt{3^{-j}} \Lambda_{n^2,k}(\sqrt{3^{-j}}\, \mathrm{dns}(T'_j)).$$

We now complete the proof of the theorem. Set $s = \sum_{i=0}^{\infty} \sqrt{3^{-i}} = 2.3660\ldots$. Then

$$\begin{aligned}
\sum_{j=0}^{\infty} \sqrt{3^{-j}} \Lambda_{n^2,k}(\sqrt{3^{-j}}\, \mathrm{dns}(T'_j)) &= s \sum_{j=0}^{\infty} \frac{\sqrt{3^{-j}}}{s} \Lambda_{n^2,k}(\sqrt{3^{-j}}\, \mathrm{dns}(T'_j)) \\
&\leq s\Lambda_{n^2,k}\left(\sum_{j=0}^{\infty} \frac{\sqrt{3^{-j}}}{s} \cdot \sqrt{3^{-j}}\, \mathrm{dns}(T'_j) \right) \\
&\leq 3\Lambda_{n^2,k}\left(\frac{s}{3} \sum_{j=0}^{\infty} \frac{\sqrt{3^{-j}}}{s} \cdot \sqrt{3^{-j}}\, \mathrm{dns}(T'_j) \right) \\
&\leq 3\Lambda_{n^2,k}(p), \qquad (6.4.15)
\end{aligned}$$

where the second step is valid by Lemma 6.24 (iii); the third step uses Corollary 6.26 with $q = s/3$; and the final step is justified by (6.4.14) and Lemma 6.24 (ii). As a result,

$$\begin{aligned}
\|T|_{\mathcal{S}}\| &\leq \sum_{j=0}^{\infty} \|T'_j|_{\mathcal{S}'}\| \\
&\leq 8C^k\, 12^{\ell-2} \binom{|I_1|}{k_1}^{1/2} \cdots \binom{|I_\ell|}{k_\ell}^{1/2} \sum_{j=0}^{\infty} \sqrt{3^{-j}} \Lambda_{n^2,k}(\sqrt{3^{-j}}\, \mathrm{dns}(T'_j)) \\
&\leq 2C^k\, 12^{\ell-1} \binom{|I_1|}{k_1}^{1/2} \cdots \binom{|I_\ell|}{k_\ell}^{1/2} \Lambda_{n^2,k}(p),
\end{aligned}$$

where the first step is valid by Proposition 2.2 and Claim 6.30, bearing in mind once again that all but finitely many of the $T'_j|_{\mathcal{S}'}$ are identically zero; the second step is a substitution from Claim 6.31; and the final step uses (6.4.15). This completes the inductive step. □

Proof of Claim 6.30. Let T' be the depth-d' decision tree given by

$$T'(v) = \begin{cases} T(v) & \text{if } v \in \{-1,1\}^{\leq d'-1}, \\ T_v|_{\binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}} & \text{if } v \in \{-1,1\}^{d'}. \end{cases}$$

This definition implies that

$$T'(v) = \begin{cases} T'_0(v) = T'_1(v) = T'_2(v) = \cdots & \text{if } v \in \{-1,1\}^{\leq d'-1}, \\ T'_0(v) + T'_1(v) + T'_2(v) + \cdots & \text{if } v \in \{-1,1\}^{d'}. \end{cases}$$

As a result,

$$\begin{aligned} T'|_{\mathcal{S}'} &= \sum_{S \in \mathcal{S}'} \sum_{v \in \{-1,1\}^{d'}} \left(\sum_{j=0}^{\infty} T'_j(v) \right) \cdot 2^{-d'} \prod_{i \in S} v_i x_{T'(v_1 v_2 \ldots v_{i-1})} \\ &= \sum_{j=0}^{\infty} \sum_{S \in \mathcal{S}'} \sum_{v \in \{-1,1\}^{d'}} T'_j(v) \cdot 2^{-d'} \prod_{i \in S} v_i x_{T'_j(v_1 v_2 \ldots v_{i-1})} \\ &= \sum_{j=0}^{\infty} T'_j|_{\mathcal{S}'}. \end{aligned} \tag{6.4.16}$$

Thus, the proof will be complete once we show that $T'|_{\mathcal{S}'} = T|_{\mathcal{S}}$.

Since $\mathcal{S}$ is the family of sets S expressible as $S = S' \cup S''$ with $S' \in \mathcal{S}'$ and $S'' \in \binom{I_\ell}{k_\ell}$, we have

$$\begin{aligned} T|_{\mathcal{S}} &= \sum_{S \in \mathcal{S}} \sum_{v \in \{-1,1\}^{d}} T(v) \cdot 2^{-d} \prod_{i \in S} v_i x_{T(v_1 v_2 \ldots v_{i-1})} \\ &= \sum_{S' \in \mathcal{S}'} \sum_{S'' \in \binom{I_\ell}{k_\ell}} \sum_{v \in \{-1,1\}^{d}} T(v) \cdot 2^{-d} \prod_{i \in S' \cup S''} v_i x_{T(v_1 v_2 \ldots v_{i-1})}. \end{aligned} \tag{6.4.17}$$

Recall that $\mathcal{S}' \subseteq \mathcal{P}(\{1, 2, \ldots, d'\})$ and $I_\ell = \{d'+1, d'+2, \ldots, d\}$. As a result, (6.4.17) yields

$$T|_{\mathcal{S}} = \sum_{S' \in \mathcal{S}'} \sum_{S'' \in \binom{I_\ell}{k_\ell}} \sum_{\substack{v' \in \{-1,1\}^{d'} \\ v'' \in \{-1,1\}^{d-d'}}} T(v'v'') \cdot 2^{-d} \prod_{i \in S'} v'_i x_{T(v'_1 v'_2 \ldots v'_{i-1})} \\ \times \prod_{i \in S''} v''_{i-d'} x_{T(v' v''_1 v''_2 \ldots v''_{i-1-d'})}.$$

A change of index now gives

$$T|_{\mathcal{S}} = \sum_{S' \in \mathcal{S}'} \sum_{S'' \in \binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}} \sum_{\substack{v' \in \{-1,1\}^{d'} \\ v'' \in \{-1,1\}^{d-d'}}} T(v'v'') \cdot 2^{-d} \prod_{i \in S'} v'_i x_{T(v'_1 v'_2 \ldots v'_{i-1})} \\ \times \prod_{i \in S''} v''_i x_{T(v' v''_1 v''_2 \ldots v''_{i-1})}.$$

Since $T(v'v'') = T_{v'}(v'')$ and $T(v'v''_1 v''_2 \ldots v''_{i-1}) = T_{v'}(v''_1 v''_2 \ldots v''_{i-1})$, we arrive at

$$T|_{\mathcal{S}} = \sum_{S' \in \mathcal{S}'} \sum_{v' \in \{-1,1\}^{d'}} 2^{-d'} \prod_{i \in S'} v'_i x_{T(v'_1 v'_2 \ldots v'_{i-1})} \\ \times \left(\sum_{S'' \in \binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}} \sum_{v'' \in \{-1,1\}^{d-d'}} T_{v'}(v'') \cdot 2^{-d+d'} \prod_{i \in S''} v''_i x_{T_{v'}(v''_1 v''_2 \ldots v''_{i-1})} \right).$$

The large parenthesized expression is by definition $T_{v'}|_{\binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}} = T'(v')$, whence

$$\begin{aligned} T|_{\mathcal{S}} &= \sum_{S' \in \mathcal{S}'} \sum_{v' \in \{-1,1\}^{d'}} T'(v') \cdot 2^{-d'} \prod_{i \in S'} v'_i x_{T(v'_1 v'_2 \ldots v'_{i-1})} \\ &= \sum_{S' \in \mathcal{S}'} \sum_{v' \in \{-1,1\}^{d'}} T'(v') \cdot 2^{-d'} \prod_{i \in S'} v'_i x_{T'(v'_1 v'_2 \ldots v'_{i-1})} \\ &= T'|_{\mathcal{S}'}. \end{aligned} \tag{6.4.18}$$

By (6.4.16) and (6.4.18), the proof is complete. □

Proof of Claim 6.31. Recall from Lemma 6.22 that $T_v|_{\binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}}$ is the k_ℓ-th homogeneous part of the function computed by the subtree T_v of T. This implies that $T'_j \in \mathcal{T}(n, d', \operatorname{dns}(T'_j), k_\ell)$. Moreover, every nonzero leaf v of T'_j has norm

$$\begin{aligned}\left\|\left| T_v|_{\binom{\{1,2,\ldots,|I_\ell|\}}{k_\ell}} \right|\right\| &\le 2C^{k_\ell} \binom{|I_\ell|}{k_\ell}^{1/2} \Lambda_{n^2,k_\ell}(\operatorname{dns}(T_v)) \\ &\le 2C^{k_\ell} \binom{|I_\ell|}{k_\ell}^{1/2} \Lambda_{n^2,k_\ell}(3^{-j}),\end{aligned}$$

where the first step applies the inductive hypothesis to the tree T_v of depth $|I_\ell|$, and the second step is legitimate by the monotonicity of Λ_{n^2,k_ℓ} (Lemma 6.24). Now Lemma 6.23 gives, for each $i = 1, 2, \ldots, \binom{n}{k_\ell}$, a real number $0 \le p_i \le 1$ and a decision tree $U_{j,i} \in \mathcal{T}^*(n, d', p_i, 0)$ such that

$$\operatorname{dns}(T'_j) = \sum_{i=1}^{\binom{n}{k_\ell}} p_i, \tag{6.4.19}$$

$$\| T'_j|_{\mathcal{S}'} \| \le 2C^{k_\ell} \binom{|I_\ell|}{k_\ell}^{1/2} \Lambda_{n^2,k_\ell}(3^{-j}) \sum_{i=1}^{\binom{n}{k_\ell}} \| U_{j,i}|_{\mathcal{S}'} \|. \tag{6.4.20}$$

Applying the inductive hypothesis to each $U_{j,i}|_{\mathcal{S}'}$ gives

$$\sum_{i=1}^{\binom{n}{k_\ell}} \| U_{j,i}|_{\mathcal{S}'} \| \le 2C^{k-k_\ell}\, 12^{\ell-2} \sqrt{\binom{|I_1|}{k_1} \cdots \binom{|I_{\ell-1}|}{k_{\ell-1}}} \sum_{i=1}^{\binom{n}{k_\ell}} \Lambda_{n^2,k-k_\ell}(p_i). \tag{6.4.21}$$

The final summation can be bounded via

$$\begin{aligned}\sum_{i=1}^{\binom{n}{k_\ell}} \Lambda_{n^2,k-k_\ell}(p_i) &\leq \binom{n}{k_\ell} \cdot \Lambda_{n^2,k-k_\ell}\left(\binom{n}{k_\ell}^{-1} \sum_{i=1}^{\binom{n}{k_\ell}} p_i\right) \\ &= n^2 \cdot \frac{1}{n^2}\binom{n}{k_\ell} \cdot \Lambda_{n^2,k-k_\ell}\left(\binom{n}{k_\ell}^{-1} \operatorname{dns}(T'_j)\right) \\ &\leq n^2 \Lambda_{n^2,k-k_\ell}\left(\frac{\operatorname{dns}(T'_j)}{n^2}\right), \end{aligned} \tag{6.4.22}$$

where the first step is valid by Lemma 6.24 (iii); the second step is a substitution from (6.4.19); and the third step uses $k_\ell \leq 2$ along with Corollary 6.26. Now

$$\begin{aligned}\|T'_j|_{\mathcal{S}'}\| &\leq 4C^k\, 12^{\ell-2}\sqrt{\binom{|I_1|}{k_1}\cdots\binom{|I_\ell|}{k_\ell}} \cdot \Lambda_{n^2,k_\ell}(3^{-j}) \cdot n^2 \Lambda_{n^2,k-k_\ell}\left(\frac{\operatorname{dns}(T'_j)}{n^2}\right) \\ &\leq 8C^k\, 12^{\ell-2}\sqrt{\binom{|I_1|}{k_1}\cdots\binom{|I_\ell|}{k_\ell}} \cdot \frac{\Lambda_{n^2,k_\ell}(\sqrt{3^{-j}})}{\sqrt{3^j}} \cdot n^2 \Lambda_{n^2,k-k_\ell}\left(\frac{\operatorname{dns}(T'_j)}{n^2}\right) \\ &\leq 8C^k\, 12^{\ell-2}\sqrt{\binom{|I_1|}{k_1}\cdots\binom{|I_\ell|}{k_\ell}} \cdot \sqrt{3^{-j}}\Lambda_{n^2,k}(\sqrt{3^{-j}}\operatorname{dns}(T'_j)), \end{aligned}$$

where the first step combines (6.4.20)–(6.4.22); the second step uses $k_\ell \leq 2$ and Corollary 6.28; and the third step applies Corollary 6.27. □

6.4.4. Generalization to elementary families. En route to our main result on the Fourier spectrum of decision trees, we now generalize Theorem 6.29 to arbitrary elementary families $\mathcal{E}$.

THEOREM 6.32. *Let $T \in \mathcal{T}^*(n,d,p,0)$ be given, for some $0 \leq p \leq 1$ and integers $n, d \geq 1$. Let k be an integer with $1 \leq k \leq d$. Then every elementary family $\mathcal{E} \subseteq \mathcal{P}_{d,k}$ satisfies*

$$\|T|_{\mathcal{E}}\| \leq (12C)^k \Lambda_{n^2,k}(p)\sqrt{|\mathcal{E}|}, \tag{6.4.23}$$

where $C \geq 1$ is the absolute constant from Theorem 6.11*.*

Proof. If $\mathcal{E} = \varnothing$, then $T|_{\mathcal{E}} \equiv 0$ and the claimed upper bound holds trivially. In the complementary case of nonempty $\mathcal{E}$, let ℓ be the minimum positive integer such that

$$\mathcal{E} = \binom{I_1}{k_1} * \binom{I_2}{k_2} * \cdots * \binom{I_\ell}{k_\ell} \tag{6.4.24}$$

for some pairwise disjoint integer intervals $I_1, I_2, \ldots, I_\ell$ and some $k_1, k_2, \ldots, k_\ell \in \{0, 1, 2\}$. Since $\mathcal{E} \neq \varnothing$, Proposition 6.14 (i) implies that $\binom{I_j}{k_j} \neq \varnothing$ for all j and therefore

$$|I_j| \geq k_j, \qquad\qquad j = 1, 2, \ldots, \ell. \tag{6.4.25}$$

The reader will recall from the definition of the $*$ operator that

$$|\mathcal{E}| = \prod_{j=1}^{\ell} \binom{|I_j|}{k_j}, \tag{6.4.26}$$

$$k = \sum_{j=1}^{\ell} k_j. \tag{6.4.27}$$

Since we chose a representation (6.4.24) with the minimum ℓ, Proposition 6.14 (ii) additionally implies that $\binom{I_j}{k_j} \neq \{\varnothing\}$ for all j, forcing

$$k_j \in \{1, 2\}, \qquad\qquad j = 1, 2, \ldots, \ell. \tag{6.4.28}$$

The previous two equations yield

$$\ell \leq k. \tag{6.4.29}$$

It follows from (6.4.25) and (6.4.28) that each I_j is a nonempty subset of $\{1, 2, \ldots, d\}$. Furthermore, by renumbering the intervals if necessary, we may assume that $I_1 < I_2 < \cdots < I_\ell$. We abbreviate $I = I_1 \cup I_2 \cup \cdots \cup I_\ell$ and $\overline{I} = \{1, 2, \ldots, d\} \setminus I$.

It is obvious that every string $v \in \{-1,1\}^d$ is uniquely determined by its substrings $v|_I$ and $v|_{\overline{I}}$. Similarly, for every $i \in I$, the prefix $v_1 v_2 \dots v_{i-1}$ is uniquely determined by the substrings $(v_1 v_2 \dots v_{i-1})|_I$ and $v|_{\overline{I}}$. This means in particular that

$$T(v) = U_{v|_{\overline{I}}}(v|_I), \qquad v \in \{-1,1\}^d \tag{6.4.30}$$

$$T(v_1 v_2 \dots v_{i-1}) = U_{v|_{\overline{I}}}((v_1 v_2 \dots v_{i-1})|_I), \qquad v \in \{-1,1\}^d, \;\; i \in I, \tag{6.4.31}$$

where $\{U_w \colon w \in \{-1,1\}^{|\overline{I}|}\}$ is a suitable collection of decision trees of depth I. By definition,

$$U_w \in \mathcal{T}^*(n, |I|, \operatorname{dns}(U_w), 0), \qquad w \in \{-1,1\}^{|\overline{I}|}. \tag{6.4.32}$$

Moreover, the densities of the U_w are related in a natural way to the density of T. Indeed, considering a uniformly random string $v \in \{-1,1\}^d$ in (6.4.30) gives $\mathbf{P}[T(v) \neq 0] = \mathbf{P}[U_{v|_{\overline{I}}}(v|_I) \neq 0]$, which is equivalent to

$$\operatorname{dns}(T) = \mathbf{E} \operatorname{dns}(U_{v|_{\overline{I}}}). \tag{6.4.33}$$

In what follows, all expectations are with respect to uniformly random $v \in \{-1,1\}^d$. We have:

$$\begin{aligned} T|_{\mathcal{E}} &= \mathbf{E}\left[\sum_{S \in \mathcal{E}} T(v) \prod_{i \in S} v_i x_{T(v_1 v_2 \dots v_{i-1})}\right] \\ &= \mathbf{E}\left[\sum_{S_1 \in \binom{I_1}{k_1}} \cdots \sum_{S_\ell \in \binom{I_\ell}{k_\ell}} T(v) \prod_{j=1}^{\ell} \prod_{i \in S_j} v_i x_{T(v_1 v_2 \dots v_{i-1})}\right] \\ &= \mathbf{E}\left[\sum_{S_1 \in \binom{I_1}{k_1}} \cdots \sum_{S_\ell \in \binom{I_\ell}{k_\ell}} U_{v|_{\overline{I}}}(v|_I) \prod_{j=1}^{\ell} \prod_{i \in S_j} v_i x_{U_{v|_{\overline{I}}}((v_1 v_2 \dots v_{i-1})|_I)}\right], \end{aligned}$$

where the last step uses (6.4.30) and (6.4.31). It remains to shift the indexing variable i. For this, let $I'_1 < I'_2 < \cdots < I'_\ell$ denote the integer intervals that form a partition of

$\{1, 2, \ldots, |I|\}$ and satisfy $|I'_j| = |I_j|$ for all j. Now the previous equation for $T|_{\mathcal{E}}$ can be restated as

$$\begin{aligned} T|_{\mathcal{E}} &= \mathbf{E}\left[\sum_{S_1 \in \binom{I'_1}{k_1}} \cdots \sum_{S_\ell \in \binom{I'_\ell}{k_\ell}} U_{v|_{\overline{I}}}(v|_I) \prod_{j=1}^{\ell} \prod_{i \in S_j} (v|_I)_i \cdot x_{U_{v|_{\overline{I}}}((v|_I)_{\leq i-1})}\right] \\ &= \mathbf{E}\left[U_{v|_{\overline{I}}}|_{\binom{I'_1}{k_1} * \cdots * \binom{I'_\ell}{k_\ell}}\right]. \end{aligned} \tag{6.4.34}$$

As a result,

$$\begin{aligned} \|T|_{\mathcal{E}}\| &\leq \mathbf{E}\left\|U_{v|_{\overline{I}}}|_{\binom{I'_1}{k_1} * \cdots * \binom{I'_\ell}{k_\ell}}\right\| \\ &\leq \mathbf{E}\left[2C^k\, 12^{\ell-1} \Lambda_{n^2,k}(\operatorname{dns}(U_{v|_{\overline{I}}})) \prod_{i=1}^{\ell} \binom{|I'_i|}{k_i}^{1/2}\right] \\ &= 2C^k\, 12^{\ell-1}\, \mathbf{E}\left[\Lambda_{n^2,k}(\operatorname{dns}(U_{v|_{\overline{I}}})) \prod_{i=1}^{\ell} \binom{|I_i|}{k_i}^{1/2}\right] \\ &= 2C^k\, 12^{\ell-1} \sqrt{|\mathcal{E}|}\ \mathbf{E}\left[\Lambda_{n^2,k}(\operatorname{dns}(U_{v|_{\overline{I}}}))\right] \\ &\leq 2C^k\, 12^{\ell-1} \sqrt{|\mathcal{E}|}\, \Lambda_{n^2,k}(\mathbf{E}\operatorname{dns}(U_{v|_{\overline{I}}})) \\ &\leq (12C)^k \sqrt{|\mathcal{E}|}\, \Lambda_{n^2,k}(\operatorname{dns}(T)), \end{aligned}$$

where the first step applies Proposition 2.2 to (6.4.34); the second step is justified by (6.4.32) and Theorem 6.29; the fourth step is a substitution from (6.4.26); the fifth step is legitimate by Lemma 6.24 (iii); and the final step uses (6.4.29) and (6.4.33). Since T has density p by hypothesis, the proof is complete. □

6.4.5. Fourier spectrum at first two levels. In this subsection, we present a proof of Theorem 6.11 for completeness. We need the following bound on the cumulative distribution function of normal distribution.

Theorem 6.33 (CDF of normal distribution [**48**]). *Let Φ be the CDF of normal distribution, i.e.,*

$$\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{x} \exp\left(-\frac{t^2}{2}\right) dt.$$

For $x > 0$,

$$1 - \Phi(x) \leq \frac{\exp(-x^2/2)}{x\sqrt{2/\pi}}.$$

With this tool at our disposal, we proceed to our proof of the first bound in Theorem 6.11.

Theorem 6.34 (Tal). *For some absolute constant $C \geq 1$, any decision tree $T : \{-1,1\}^n \to \{0,1\}$ of depth d and density $\operatorname{dns}(T) = p$, we have*

$$\| L_1 T \| \leq \binom{d}{1}^{1/2} Cp\sqrt{\ln\frac{e}{p}}. \tag{6.4.35}$$

Proof. We can assume $p \leq 0.5$, since

$$\| L_1 T \| = \| L_1 (1 - T) \|,$$

and the bound (6.4.35) is increasing in $p \in [0,1]$. For arbitrary $\sigma \in \{-1,1\}^n$, and any path $v \in \{-1,1\}^d$, abbreviate

$$(\sigma, v)_T = \sum_{i=1}^{d} \sigma_{T(v_{<i})} v_i.$$

Observe that

$$\| L_1 T \| = \max_{\sigma} \mathop{\mathbf{E}}_{v \in \{-1,1\}^d} [T(v) \cdot (\sigma, v)_T]. \tag{6.4.36}$$

The above equation holds because $\mathbf{E}_{v\in\{-1,1\}^d}[T(v)\sum_{i=1}^{d} v_i]$ is simply $\sum_{i=1}^{n} \hat{T}(i)$. Here σ_i works as guessing $\operatorname{sgn} \hat{T}(i)$.

Now fix σ. Let $v \in \{-1,1\}^d$ be a uniformly random root-to-leaf path in T. Define random variable $X_i = \sigma_{T(v_{<i})} \cdot v_i$ for $i \in [d]$. Note that the partial sums

$$\sum_{i=1}^{j} X_i \qquad\qquad j = 0, 1, 2, \ldots, d$$

form a martingale. Let $\mathcal{P} \subseteq \{0,1\}^d$ denote any set of root-to-leaf paths such that

(i) $\Pr[\mathcal{P}] = p$, i.e., $|\mathcal{P}| = p2^d$,

(ii) $(\sigma, v)_T \geq (\sigma, v')_T$, $\forall\, v \in \mathcal{P}, v' \notin \mathcal{P}$.

Since $p \leq 1/2$, by symmetry we can deduce that $(\sigma, v)_T \geq 0$ for any $v \in \mathcal{P}$. For $t \geq 0$, let $\mathcal{P}_t = \{v : (\sigma, v)_T \geq t\}$. By Azuma's inequality,

$$\Pr_{v\in\{-1,1\}^d}[v \in \mathcal{P}_t] = \Pr_{v\in\{-1,1\}^d}\left[\sum_{i=1}^{d} X_i \geq t\right] \leq \exp\left(-\frac{t^2}{2d}\right). \qquad (6.4.37)$$

Set $\theta = \left\lceil \sqrt{2d \ln \frac{1}{p}} \right\rceil$. The choice of θ guarantees that $|\mathcal{P}_\theta| \leq p2^d$. Therefore, it holds that $\mathcal{P}_\theta \subseteq \mathcal{P}$ by our definition of $\mathcal{P}_\theta$ and $\mathcal{P}$. Now we bound $\mathbf{E}_v[T(v) \cdot (\sigma, v)_T]$ as follows

$$\begin{aligned}
\mathop{\mathbf{E}}_{v\in\{-1,1\}^d}[T(v) \cdot (\sigma, v)_T] &\leq 2^{-d} \sum_{v\in\mathcal{P}} (\sigma, v)_T \\
&\leq 2^{-d} \sum_{v\in\mathcal{P}\setminus\mathcal{P}_\theta} \theta + 2^{-d} \sum_{v\in\mathcal{P}_\theta} (\sigma, v)_T \\
&\leq \Pr[\mathcal{P} \setminus \mathcal{P}_\theta]\theta + \Pr[\mathcal{P}_\theta]\theta + \sum_{t=\theta+1}^{d} \Pr[\mathcal{P}_t] \\
&\leq p\theta + \sum_{t=\theta+1}^{d} \exp\left(-\frac{t^2}{2d}\right) \\
&\leq p\theta + \int_{\theta}^{\infty} \exp\left(-\frac{t^2}{2d}\right) dt \\
&\leq p\left\lceil \sqrt{2d \ln \frac{1}{p}} \right\rceil + \frac{1}{\sqrt{8}} \cdot p\sqrt{d},
\end{aligned}$$

where the first equality holds because $(\sigma, v)_T \geq 0$ for any $v \in \mathcal{P}$; the second step holds because $(\sigma, v)_T < \theta$ for $v \in \mathcal{P} \setminus \mathcal{P}_\theta$; the forth step is by (6.4.37); and last step follows Theorem 6.33. In view of the (6.4.36), we are done. □

Next we move on to our proof of the second bound in Theorem 6.11.

THEOREM 6.35 (Tal). *For some absolute constant C, any decision tree $T : \{-1,1\}^n \to \{0,1\}$ of depth d and density* $\mathrm{dns}(T) = p$, *we have*

$$\left\| L_2 T \right\| \leq \binom{d}{2}^{1/2} C^2 p \sqrt{\ln \frac{e}{p}} \sqrt{\ln \frac{en}{p}}. \tag{6.4.38}$$

Proof. For each $k \in [n]$ define a decision tree T_k, which is formed by cutting the tree T by the nodes labeled by x_k. Precisely, for any node $v \in \{-1,1\}^{<d}$, consider the following cases:

(i) $T(v) = k$. Then we set $T_k(v) = \hat{T}_v(k)$. Now v will be a leaf of T_k.

(ii) For the remaining case, set $T_k(v) = T(v)$.

All the leaves not set in the above cases are labeled by 0. Observe that for $v \in \{-1,1\}^{<d}$,

$$\hat{T}_v(T(v)) = \mathop{\mathbf{E}}_{\substack{u \in \{-1,1\}^d: \\ u \succ v}} [T(u)]. \tag{6.4.39}$$

The above process naturally gives two types of leaves A_k, B_k. A_k contains leaves v such that $|v| < d$, and these leaves are labeled by some real number which is a multiple of $2^{-(d-|v|)}$. B_k contains the leaves all labeled by 0. Let

$$p_k = \sum_{v \in A_k} 2^{-|v|} |T_k(v)|.$$

Let c be the constant in Theorem 6.34, the following three equations relate the trees constructed above and $\|\!| L_2\, T |\!\|$:

$$L_2\, T = \sum_{k=1}^{n} x_k \cdot L_1\, T_k, \tag{6.4.40}$$

$$\sum_{k=1}^{n} p_k \leq \binom{d}{1}^{1/2} cp \sqrt{\ln \frac{e}{p}}, \tag{6.4.41}$$

$$\|\!| L_1\, T_k |\!\| \leq \binom{d}{1}^{1/2} cp_k \sqrt{\ln \frac{e}{p_k}}, \qquad k = 1, 2, \ldots, n. \tag{6.4.42}$$

Assuming the above three equations are true, then

$$\begin{aligned}
\|\!| L_2\, T |\!\| &= \left\|\!\left| \sum_{k=1}^{n} x_k \cdot L_1\, T_k \right|\!\right\| \\
&\leq \sum_{k=1}^{n} \|\!| L_1\, T_k |\!\| \\
&\leq \sum_{k=1}^{n} \binom{d}{1}^{1/2} cp_k \sqrt{\ln \frac{e}{p_k}} \\
&\leq \binom{d}{1}^{1/2} c \left(\sum_{k=1}^{n} p_k \right) \sqrt{\ln \frac{en}{(\sum_{k=1}^{n} p_k)}} \\
&\leq \binom{d}{1} c^2 p \sqrt{\ln \frac{e}{p}} \sqrt{\ln \frac{en}{cp\sqrt{d \ln(e/p)}}} \\
&\leq \binom{d}{2}^{2} C^2 p \sqrt{\ln \frac{e}{p}} \sqrt{\ln \frac{en}{p}},
\end{aligned}$$

where the fourth and fifth step uses the concavity and monotonicity of the regarding functions; the last step holds since $c\sqrt{d \ln(e/p)} \geq 1$. Now we verify (6.4.40)–(6.4.42).

(6.4.40). Expand $L_2\,T$,

$$\begin{aligned}
L_2\,T(x) &= \mathop{\mathbf{E}}_{v\in\{-1,1\}^d}\left[\sum_{1\le i<j\le d} T(v)v_iv_jx_{T(v_{<i})}x_{T(v_{<j})}\right] \\
&= \sum_{k=1}^{n}\mathop{\mathbf{E}}_{v\in\{-1,1\}^d}\left[T(v)\sum_{1<j\le d} v_jx_k\mathbf{I}[T(v_{<j})=k]\cdot\left(\sum_{1\le i<j} v_ix_{T(v_{<i})}\right)\right] \\
&= \sum_{k=1}^{n}\sum_{u\in A_k}\mathop{\mathbf{E}}_{v\in\{-1,1\}^d}\left[\mathbf{I}[u\preceq v]\cdot T(v)v_{|u|}x_k\sum_{1\le i<|u|} u_ix_{T(u_{<i})}\right] \\
&= \sum_{k=1}^{n}\sum_{u\in A_k}\Pr_{v\in\{-1,1\}^d}[\mathbf{I}[u\preceq v]]\cdot\hat{T}_u(k)x_k\sum_{1\le i<|u|} u_ix_{T(u_{<i})} \\
&= \sum_{i=1}^{n} x_k\cdot L_1\,T_k,
\end{aligned}$$

where the fourth step follows (6.4.39).

(6.4.41). For each internal node $v\in\{-1,1\}^{<d}$, without loss of generality, assume that $\hat{T}_v(T(v))\ge 0$. Since otherwise, we can exchange the left and right children of v in T without changing p_k. Under this assumption, we have for any $k\in[n]$,

$$p_k = |\hat{T}(k)|.$$

By Theorem 6.34,

$$\sum_{i=1} p_k \le \binom{d}{1}^{1/2} Cp\sqrt{\ln\frac{e}{p}}.$$

(6.4.42). The inequality holds if there is a full binary tree of T' of depth d and density at most p, such that $\|\!|L_1\,T'\|\!| \ge \|\!|L_1\,T_k\|\!|$, as $\|\!|L_1\,T'\|\!|$ is be bounded by Theorem 6.34. We construct such T'. Let $T'(v)=T_k(v)$ for all internal node v of T_k, and $T'(v)=0$ for $v\in B_k$. The nodes left are those in the subtrees T_v for $v\in A_k$. Note that for any $v\in A_k$, $L_1\,T_v$ is a multiple of $2^{-(d-|v|)}$. We label T'_v in a very "generous" manner. Label each node in T'_v using a fresh variable. Consequently, T' can have an enormous number of variables, but it is not a problem for us. Label $2^{(d-|v|)}|T_k(v)|$ leaves in T'_{v1} by 1 if $T_k(v) > 0$, and label $2^{(2-|v|)}|T_k(v)|$ leaves in $T'_{v\circ -1}$ by 1 if $T_k(v) < 0$. Label all other leaves by 0. The fraction of nonzero leaves in T' is exactly

$$\sum_{v \in A_k} 2^{-|v|}|T_k(v)| = p_k.$$

It is also easy to check that for all $i \in [n] \setminus \{k\}$,

$$\hat{T'}(i) = \hat{T}_k(i).$$

Therefore $\|\!|L_1\, T'\|\!| \geq \|\!|L_1\, T_k\|\!|$. $\square$

The discussions so far settled Theorem 6.11 for decision trees with range $\{0, 1\}$. We now extend the two bounds to decision trees with range $\{-1, 0, 1\}$.

Corollary 6.36 (Theorem 6.11 restated). *For some absolute constant C, any decision tree $T \in \mathcal{T}^*(n, d, p, 0)$, we have*

$$\|\!|L_1\, T\|\!| \leq \binom{d}{1}^{1/2} Cp\sqrt{\ln \frac{e}{p}},$$

$$\|\!|L_2\, T\|\!| \leq \binom{d}{2}^{1/2} C^2 p\sqrt{\ln \frac{e}{p}}\sqrt{\ln \frac{en}{p}}.$$

Proof. Let C' be the constant in Theorem 6.34. Write $T = T^+ - T^-$, where T^+ and T^- are both decision trees with range $\{0, 1\}$. The nonzero leaves of T^+ and T^- are those leaves of T labeled by 1 and -1, respectively. Let p^+, p^- be the fraction of the

1 and -1 leaves in T, respectively. Then

$$\begin{aligned}
\|\!| L_1 T |\!\| &\le \|\!| L_1 T^+ |\!\| + \|\!| L_1 T^- |\!\| \\
&\le \binom{d}{1}^{1/2} C' p^+ \sqrt{\ln \frac{e}{p^+}} + \binom{d}{1}^{1/2} C' p^- \sqrt{\ln \frac{e}{p^-}} \\
&\le \binom{d}{1}^{1/2} C' p \sqrt{\ln \frac{2e}{p}},
\end{aligned}$$

where the last step is due to the concavity of $p\sqrt{\ln(e/p)}$ for $p \in (0,1]$. A completely analogous argument works for $\|\!| L_2 T |\!\|$ using the concavity of $p\sqrt{\ln(e/p)\ln(en/p)}$ for $p \in (0,1]$. □

6.4.6. Main result. We now obtain our main result on the Fourier spectrum of decision trees by combining Theorem 6.32 with an efficient decomposition of $\mathcal{P}_{d,k}$ into elementary families (Theorem 6.20).

THEOREM 6.37. *Let $f\colon \{-1,1\}^n \to \{-1,0,1\}$ be a function computable by a decision tree of depth d. Define $p = \mathbf{P}_{x\in\{-1,1\}^n}[f(x) \ne 0]$. Then*

$$\|\!| L_k f |\!\| \le \binom{d}{k}^{1/2} (58Cc)^k \, \Lambda_{n^2,k}(p), \qquad\qquad k = 1, 2, \ldots, n,$$

where $C \ge 1$ and $c \ge 1$ are the absolute constants from Theorem 6.11 *and Lemma* 6.12, respectively.

Proof. Lemma 6.22 ensures that $L_k f = 0$ for $k > d$, so that the theorem holds vacuously in that case. We now examine the complementary possibility, $1 \le k \le d$. For some integer $N \ge 1$, Theorem 6.20 gives a partition $\mathcal{P}_{d,k} = \bigcup_{i=1}^N \mathcal{E}_i$ where $\mathcal{E}_1, \mathcal{E}_2, \ldots, \mathcal{E}_N$ are elementary families with

$$\sum_{i=1}^N |\mathcal{E}_i|^{1/2} \le (2+2\sqrt{2})^k c^k \left(\frac{d}{k}\right)^{k/2}. \tag{6.4.43}$$

Fix a decision tree T of depth d that computes f. Then Fact 6.10 shows that $T \in \mathcal{T}^*(n, d, p, 0)$. As a result,

$$\begin{aligned}
\|\!| L_k f |\!\| &= \|\!| T|_{\mathcal{P}_{d,k}} |\!\| \\
&= \left\|\!\left| \sum_{i=1}^{N} T|_{\mathcal{E}_i} \right|\!\right\| \\
&\leq \sum_{i=1}^{N} \|\!| T|_{\mathcal{E}_i} |\!\| \\
&\leq \sum_{i=1}^{N} (12C)^k \, \Lambda_{n^2,k}(p) \sqrt{|\mathcal{E}_i|} \\
&\leq \left(\frac{d}{k}\right)^{k/2} (58Cc)^k \, \Lambda_{n^2,k}(p),
\end{aligned}$$

where the first step is valid by Lemma 6.22; the second step uses Proposition 6.21; the third step uses Proposition 2.2; the fourth step applies Theorem 6.32; and the final step substitutes the upper bound from (6.4.43). In view of (2.1.1), the proof is complete. □

Maximizing over $0 \leq p \leq 1$, we establish the following clean bound conjectured by Tal [**135**].

Corollary 6.38. *Let $f\colon \{-1,1\}^n \to \{-1,0,1\}$ be a function computable by a decision tree of depth d. Then*

$$\|\!| L_k f |\!\| \leq C^k \sqrt{\binom{d}{k} (1 + \ln n)^{k-1}}, \qquad\qquad k = 1, 2, \ldots, n,$$

where $C \geq 1$ is an absolute constant.

Proof. Recall from Lemma 6.24 (ii) that $\Lambda_{n^2,k}(p) \leq \sqrt{(\ln en^2)^{k-1}}$ for all $0 \leq p \leq 1$. Now the claimed bound is immediate from Theorem 6.37 after a change of constant C. □

Corollary 6.38 settles Theorem 6.8 from the introduction. By convexity (Proposition 2.2), Corollary 6.38 holds more generally for any real function $f\colon \{-1,1\}^n \to [-1,1]$ computable by a decision tree of depth d.

6.5. Quantum versus classical query complexity

Using our newly derived bound for the Fourier spectrum of decision trees, we will now prove the main result of this chapter on quantum versus randomized query complexity.

6.5.1. Quantum and randomized query models. For a nonempty finite set X, a *partial Boolean function on* X is a mapping $X \to \{0,1,*\}$, where the output value $*$ is reserved for illegal inputs. Recall that a *randomized query algorithm of cost* d is a probability distribution on decision trees of depth at most d. For a (possibly partial) Boolean function f on the Boolean hypercube, we say that a randomized query algorithm *computes* f *with error* ε if, for every input $x \in f^{-1}(0) \cup f^{-1}(1)$, the algorithm outputs $f(x)$ with probability at least $1-\varepsilon$. Observe that in this formalism, the algorithm is allowed to exhibit arbitrary behavior on the illegal inputs, namely, those in $f^{-1}(*)$. The *randomized query complexity* $R_\varepsilon^{\mathrm{dt}}(f)$ is the minimum cost of a randomized query algorithm that computes f with error ε. The canonical setting of the error parameter is $\varepsilon = 1/3$. This choice is largely arbitrary because the error of a query algorithm can be reduced in an efficient manner by running the algorithm several times independently and outputting the majority answer. Quantitatively, the following relation follows from the Chernoff bound:

$$R_{\varepsilon}^{\mathrm{dt}}(f) \leq O\left(\frac{1}{\gamma^2}\log\frac{1}{\varepsilon}\right) \cdot R_{\frac{1}{2}-\gamma}^{\mathrm{dt}}(f) \tag{6.5.1}$$

for all $\varepsilon, \gamma \leq 1/2$.

These classical definitions carry over in the obvious way to the quantum model. Here, the cost is the worst-case number of quantum queries on any input, and a quantum algorithm is said to *compute f with error ε* if, for every input $x \in f^{-1}(0) \cup f^{-1}(1)$, the algorithm outputs $f(x)$ with probability at least $1-\varepsilon$. The *quantum query complexity* $Q_{\varepsilon}^{\mathrm{dt}}(f)$ is the minimum cost of a quantum query algorithm that computes f with error ε. For an excellent introduction to classical and quantum query complexity, we refer the reader to [**29**] and [**138**], respectively.

6.5.2. The rorrelation problem. We now formally state the problem of interest to us, Tal's *rorrelation* [**135**], which was briefly reviewed in the introduction. Let n and k be positive integers. For an orthogonal matrix $U \in \mathbb{R}^{n\times n}$, consider the multilinear polynomial $\phi_{n,k,U} \colon (\{-1,1\}^n)^k \to \mathbb{R}$ given by

$$\phi_{n,k,U}(x_1, x_2, \ldots, x_k) = \frac{1}{n}\mathbf{1}^{\mathsf{T}} D_{x_1} U D_{x_2} U D_{x_3} U \cdots U D_{x_k} \mathbf{1}, \tag{6.5.2}$$

where $\mathbf{1}$ denotes the all-ones vector and D_{x_i} denotes the diagonal matrix with vector x_i on the diagonal. In what follows, we treat the sets $(\{-1,1\}^n)^k$ and $\{-1,1\}^{n\times k}$ interchangeably, thereby interpreting the input to $\phi_{n,k,U}$ as an $n \times k$ sign matrix. Let

$\|\cdot\|_2$ denote the Euclidean norm. Then for all $x_1, x_2, \ldots, x_k \in \{-1,1\}^n$, we have

$$\begin{aligned} |\phi_{n,k,U}(x_1, x_2, \ldots, x_k)| &= \frac{1}{n}\langle \mathbf{1}, D_{x_1} U D_{x_2} U D_{x_3} U \cdots U D_{x_k} \mathbf{1}\rangle \\ &\le \frac{1}{n}\|\mathbf{1}\|_2 \; \|D_{x_1} U D_{x_2} U D_{x_3} U \cdots U D_{x_k} \mathbf{1}\|_2 \\ &= \frac{1}{n}\|\mathbf{1}\|_2 \; \|\mathbf{1}\|_2 \\ &= 1, \end{aligned} \tag{6.5.3}$$

where the second step applies the Cauchy–Schwarz inequality, and the third step is valid because each of the matrices involved preserves the Euclidean norm. In particular, the multivariate polynomial $\phi_{n,k,U}$ ranges in $[-1,1]$ for all inputs. Generalizing the forrelation problem of Aaronson and Ambainis [**2**], Tal [**135**] considered the partial Boolean function $f_{n,k,U}\colon \{-1,1\}^{n\times k} \to \{0,1,*\}$ given by

$$f_{n,k,U}(x) = \begin{cases} 1 & \text{if } \phi_{n,k,U}(x) \ge 2^{-k}, \\ 0 & \text{if } |\phi_{n,k,U}(x)| \le 2^{-k-1}, \\ * & \text{otherwise.} \end{cases}$$

Aaronson and Ambainis [**2**] showed that there is a quantum algorithm with $\lceil k/2 \rceil$ queries whose acceptance probability on input $x \in \{-1,1\}^{n\times k}$ is $(\phi_{n,k,H}(x)+1)/2$, where H is the Hadamard transform matrix. Their analysis generalizes to any orthogonal matrix in place of H, to the following effect.

FACT 6.39 (Tal [**135**, Claim 3.1]). *Let n and k be positive integers, where n is a power of 2. Let U be an arbitrary orthogonal matrix. Then there is a quantum query algorithm with $\lceil k/2 \rceil$ queries whose acceptance probability on input $x \in \{-1,1\}^{n\times k}$ equals $(\phi_{n,k,U}(x)+1)/2$.*

Corollary 6.40. *Let n and k be positive integers, where n is a power of 2. Let U be an arbitrary orthogonal matrix. Then*

$$Q^{\mathrm{dt}}_{\frac{1}{2}-\frac{1}{2^{k+4}}}(f_{n,k,U}) \leq \left\lceil \frac{k}{2} \right\rceil. \tag{6.5.4}$$

In particular,

$$Q^{\mathrm{dt}}_{1/3}(f_{n,k,U}) \leq O(k4^k). \tag{6.5.5}$$

Proof. On input x, the query algorithm for (6.5.4) is as follows: with probability p, run the algorithm of Fact 6.39 and output the resulting answer; with complementary probability $1-p$, output “no” regardless of x. By design, the proposed solution has query cost at most $\lceil k/2 \rceil$ and accepts x with probability exactly

$$p \cdot \frac{\phi_{n,k,U}(x)+1}{2}.$$

We want this quantity to be at most $\frac{1}{2} - 2^{-k-4}$ if $\phi_{n,k,U}(x) \leq 2^{-k-1}$, and at least $\frac{1}{2} + 2^{-k-4}$ if $\phi_{n,k,U}(x) \geq 2^{-k}$. These requirements are both met for $p = (1+\frac{3}{2^{k+2}})^{-1}$. In summary, $f_{n,k,U}$ has a query algorithm with error at most $\frac{1}{2} - 2^{-k-4}$ and query cost $\lceil k/2 \rceil$. To reduce the error to 1/3, run this algorithm independently $\Theta(4^k)$ times and output the majority answer; cf. (6.5.1). □

Corollary 6.40 shows that the rorrelation problem has small quantum query complexity. By contrast, we will show that its randomized complexity is essentially the maximum possible. Specifically, we will prove an optimal, near-linear lower bound on the randomized query complexity of rorrelation by combining Tal's work [**135**] with our near-optimal bounds for the Fourier spectrum of decision trees.

In what follows, let $\mathcal{U}_{n,k}$ denote the uniform probability distribution on $\{-1,1\}^{n\times k}$. Applying Parseval's identity to the multilinear polynomial $\phi_{n,k,U}$ gives:

FACT 6.41 (Tal [**135**, Claim 4.4]). $\mathbf{E}_{x\sim\mathcal{U}_{n,k}}[\phi_{n,k,U}(x)^2] = 1/n$.

The other result from [**135**] that we will need is as follows.

FACT 6.42 (Tal [**135**, Lemmas 5.6, 5.7, and Claim 4.1]). *Let n and k be positive integers. Let $U \in \mathbb{R}^{n\times n}$ be a uniformly random orthogonal matrix. Then with probability $1 - o(1)$, there exists a probability distribution $\mathcal{D}_{n,k,U}$ on $\{-1,1\}^{n\times k}$ such that:*

$$\mathop{\mathbf{E}}_{x\sim\mathcal{D}_{n,k,U}} \phi_{n,k,U}(x) \geq \left(\frac{2}{\pi}\right)^{k-1}, \tag{6.5.6}$$

$$\mathop{\mathbf{E}}_{x\sim\mathcal{D}_{n,k,U}} \prod_{(i,j)\in S} x_{i,j} = 0, \qquad |S| = 1, 2, \ldots, k-1, \tag{6.5.7}$$

$$\left|\mathop{\mathbf{E}}_{x\sim\mathcal{D}_{n,k,U}} \prod_{(i,j)\in S} x_{i,j}\right| \leq \left(\frac{c|S|\log n}{n}\right)^{\frac{|S|}{2}\cdot\frac{k-1}{k}}, \qquad |S| = k, k+1, \ldots, nk, \tag{6.5.8}$$

where $c \geq 1$ is an absolute constant independent of n, k, U.

6.5.3. The quantum-classical separation. In this section, we derive our lower bound on the randomized query complexity of the rorrelation problem by combining Tal's Facts 6.41 and 6.42 with our main result on decision trees (Corollary 6.38). The technical centerpiece of this derivation is the following "indistinguishability" lemma, which is a polynomial improvement on the analogous calculation by Tal [**135**, Theorem 5.8] that used weaker Fourier bounds for decision trees.

LEMMA 6.43. *Let n and k be positive integers. Let $U \in \mathbb{R}^{n\times n}$ be a uniformly random orthogonal matrix. Then with probability $1 - o(1)$, every function $g\colon \{-1,1\}^{n\times k} \to \{0,1\}$ obeys*

$$\left|\mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} g - \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} g\right| \leq \left(cd\cdot\frac{\log^{2-\frac{1}{k}}(n+k)}{n^{1-\frac{1}{k}}}\right)^{k/2}, \tag{6.5.9}$$

where $\mathcal{D}_{n,k,U}$ is as defined in Fact 6.42; *d is the minimum depth of a decision tree that computes g; and $c \geq 1$ is an absolute constant independent of n, k, U, g.*

Proof. Fact 6.42 guarantees that with probability $1 - o(1)$, there is a probability distribution $\mathcal{D}_{n,k,U}$ on $\{-1,1\}^{n\times k}$ that obeys (6.5.6)–(6.5.8). Conditioned on this event, we will prove (6.5.9). To start with, fix g and write out the Fourier expansion

$$\begin{aligned} g(x) &= \sum_{S\subseteq\{1,2,\ldots,n\}\times\{1,2,\ldots,k\}} \hat{g}(S) \prod_{(i,j)\in S} x_{i,j} \\ &= \sum_{\ell=0}^{nk} \sum_{|S|=\ell} \hat{g}(S) \prod_{(i,j)\in S} x_{i,j}. \end{aligned}$$

Then

$$\begin{aligned} \left| \mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} g - \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} g \right| &\leq \sum_{\ell=0}^{nk} \sum_{|S|=\ell} |\hat{g}(S)| \left| \mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} \prod_{(i,j)\in S} x_{i,j} - \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} \prod_{(i,j)\in S} x_{i,j} \right| \\ &\leq \sum_{\ell=1}^{nk} \sum_{|S|=\ell} |\hat{g}(S)| \left| \mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} \prod_{(i,j)\in S} x_{i,j} - \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} \prod_{(i,j)\in S} x_{i,j} \right| \\ &\leq \sum_{\ell=k}^{nk} \sum_{|S|=\ell} |\hat{g}(S)| \left| \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} \prod_{(i,j)\in S} x_{i,j} \right|, \end{aligned}$$

where the first step uses the triangle inequality; the second step is justified by $\mathbf{E}_{\mathcal{U}_{n,k}} 1 = \mathbf{E}_{\mathcal{D}_{n,k,U}} 1 = 1$; and the third step is valid due to (6.5.7) and the identity $\mathbf{E}_{\mathcal{U}_{n,k}} \prod_{(i,j)\in S} x_{i,j} = 0$ for nonempty S. Let d be the minimum depth of a decision tree that computes g. Applying (6.5.8) then Corollary 6.38, we conclude that

$$\left| \mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} g - \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} g \right| \leq \sum_{\ell=k}^{nk} c_1^{\ell} \sqrt{\binom{d}{\ell} (1 + \ln nk)^{\ell-1}} \left(\frac{c_2 \ell \log n}{n} \right)^{\frac{\ell}{2} \cdot \frac{k-1}{k}},$$

where $c_1 \geq 1$ and $c_2 \geq 1$ are the absolute constants in Corollary 6.38 and Fact 6.42. In view of (2.1.1), this gives

$$\left|\mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} g - \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} g\right| \leq \sum_{\ell=k}^{\infty} \left(c_1^2 \cdot \frac{ed}{\ell} \cdot (1 + \ln nk)^{\frac{\ell-1}{\ell}} \cdot \left(\frac{c_2 \ell \log n}{n}\right)^{\frac{k-1}{k}} \right)^{\frac{\ell}{2}}$$

$$\leq \sum_{\ell=k}^{\infty} \left(c_1^2 \cdot ed \cdot (1 + \ln nk) \cdot \left(\frac{c_2 \log n}{n}\right)^{\frac{k-1}{k}} \right)^{\frac{\ell}{2}}$$

$$\leq \sum_{\ell=k}^{\infty} \left(\frac{cd}{4} \cdot \frac{\log^{2-\frac{1}{k}}(n+k)}{n^{1-\frac{1}{k}}} \right)^{\frac{\ell}{2}},$$

where $c \geq 1$ in the last step is a sufficiently large absolute constant. This settles (6.5.9) in the case when $cd \log^{(2k-1)/k}(n+k) \leq n^{(k-1)/k}$. In the complementary case, (6.5.9) follows from the trivial bound $|\mathbf{E}_{\mathcal{U}_{n,k}} g - \mathbf{E}_{\mathcal{D}_{n,k,U}} g| \leq 1$. □

We have reached the main result of this section, an essentially tight lower bound on the randomized query complexity of the k-fold rorrelation problem.

THEOREM 6.44. *Let n and k be positive integers, with $k \leq \frac{1}{3}\log n - 1$. Let $U \in \mathbb{R}^{n\times n}$ be a uniformly random orthogonal matrix. Then with probability $1 - o(1)$,*

$$R^{\mathrm{dt}}_{1/2^{k+1}}(f_{n,k,U}) = \Omega\left(\frac{n^{1-\frac{1}{k}}}{(\log n)^{2-\frac{1}{k}}}\right) \tag{6.5.10}$$

and in particular

$$R^{\mathrm{dt}}_{\frac{1}{2}-\gamma}(f_{n,k,U}) = \Omega\left(\frac{\gamma^2}{k} \cdot \frac{n^{1-\frac{1}{k}}}{(\log n)^{2-\frac{1}{k}}}\right), \qquad 0 \leq \gamma \leq \frac{1}{2}. \tag{6.5.11}$$

Proof. We will prove the lower bound for every U that satisfies (6.5.6) and (6.5.9), which happens with probability $1 - o(1)$ by Fact 6.42 and Lemma 6.43. To begin with,

$$\begin{aligned}\mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[f_{n,k,U}(x) \neq 0] &= \mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[|\phi_{n,k,U}(x)| > 2^{-k-1}] \\ &\leq 4^{k+1} \mathop{\mathbf{E}}_{\mathcal{U}_{n,k}}[\phi_{n,k,U}(x)^2] \\ &\leq \frac{4^{k+1}}{n} \\ &\leq \frac{1}{2^{k+1}}, \end{aligned} \tag{6.5.12}$$

where the last three steps use Markov's inequality, Fact 6.41, and $k \leq \frac{1}{3}\log n - 1$, respectively. Also,

$$\begin{aligned}\left(\frac{2}{\pi}\right)^{k-1} &\leq \mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} \phi_{n,k,U}(x) \\ &\leq 2^{-k} \mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[\phi_{n,k,U}(x) < 2^{-k}] + \mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[\phi_{n,k,U}(x) \geq 2^{-k}] \\ &= 2^{-k}(1 - \mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x) = 1]) + \mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x) = 1] \\ &= 2^{-k} + (1 - 2^{-k}) \mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x) = 1], \end{aligned}$$

where the first and second steps are justified by (6.5.6) and (6.5.3), respectively. The last equation shows that

$$\begin{aligned}\mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x) = 1] &\geq \left(\frac{2}{\pi}\right)^{k-1} - 2^{-k} \\ &\geq 2^{-k}. \end{aligned} \tag{6.5.13}$$

Now fix arbitrary parameters $d \geq 1$ and $0 \leq \varepsilon \leq 1/2$, and consider a randomized query algorithm of cost d that computes $f_{n,k,U}$ with error at most ε. Then the algorithm's acceptance probability on given input x is $\mathbf{E}_r\, g_r(x)$, where r denotes a random string and each $g_r\colon \{-1,1\}^{n\times k} \to \{0,1\}$ is computable by a decision tree of depth at most

d. Since the error is at most ε, we have

$$\mathop{\mathbf{P}}_{r}[f_{n,k,U}(x)=0,\ g_r(x)=1]+\mathop{\mathbf{P}}_{r}[f_{n,k,U}(x)=1,\ g_r(x)=0]\le\varepsilon \tag{6.5.14}$$

for every $x\in\{-1,1\}^{n\times k}$. We thus obtain the two inequalities

$$\mathop{\mathbf{E}}_{r}\mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[f_{n,k,U}(x)=0,\ g_r(x)=1]\le\varepsilon, \tag{6.5.15}$$

$$\mathop{\mathbf{E}}_{r}\mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x)=1,\ g_r(x)=0]\le\varepsilon, \tag{6.5.16}$$

by passing to expectations in (6.5.14) with respect to $x\sim\mathcal{U}_{n,k}$ and $x\sim\mathcal{D}_{n,k,U}$, respectively. On the other hand, (6.5.9) and $k=O(\log n)$ imply

$$\mathop{\mathbf{E}}_{r}\left|\mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}}g_r-\mathop{\mathbf{E}}_{\mathcal{U}_{n,k}}g_r\right|\le\left(c'd\cdot\frac{(\log n)^{2-\frac{1}{k}}}{n^{1-\frac{1}{k}}}\right)^{\frac{k}{2}} \tag{6.5.17}$$

for some absolute constant $c'\ge1$.

We now have all the ingredients to complete the proof. For each r, we have

$$\begin{aligned}\mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}}g_r&=\mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[g_r(x)=1]\\&\ge\mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x)=1]-\mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x)=1,\ g_r(x)=0]\\&\ge2^{-k}-\mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x)=1,\ g_r(x)=0],\end{aligned} \tag{6.5.18}$$

where the last step uses (6.5.13). Similarly,

$$\begin{aligned}\mathop{\mathbf{E}}_{\mathcal{U}_{n,k}}g_r&=\mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[g_r(x)=1]\\&\le\mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[f_{n,k,U}(x)\ne0]+\mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[f_{n,k,U}(x)=0,\ g_r(x)=1]\\&\le2^{-k-1}+\mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[f_{n,k,U}(x)=0,\ g_r(x)=1],\end{aligned} \tag{6.5.19}$$

where the last step uses (6.5.12). Passing to expectations in (6.5.18) and (6.5.19) with respect to r gives

$$\mathop{\mathbf{E}}_{r}\left[\mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} g_r - \mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} g_r\right] \geq 2^{-k-1} - \mathop{\mathbf{E}}_{r}\mathop{\mathbf{P}}_{\mathcal{D}_{n,k,U}}[f_{n,k,U}(x)=1,\ g_r(x)=0]$$
$$- \mathop{\mathbf{E}}_{r}\mathop{\mathbf{P}}_{\mathcal{U}_{n,k}}[f_{n,k,U}(x)=0,\ g_r(x)=1],$$

which in view of (6.5.15) and (6.5.16) simplifies to

$$\mathop{\mathbf{E}}_{r}\left[\mathop{\mathbf{E}}_{\mathcal{D}_{n,k,U}} g_r - \mathop{\mathbf{E}}_{\mathcal{U}_{n,k}} g_r\right] \geq 2^{-k-1} - 2\varepsilon.$$

Comparing this lower bound with (6.5.17), we arrive at

$$\left(c'd \cdot \frac{(\log n)^{2-\frac{1}{k}}}{n^{1-\frac{1}{k}}}\right)^{\frac{k}{2}} \geq 2^{-k-1} - 2\varepsilon.$$

Taking $\varepsilon = 2^{-k-3}$ and solving for d, we find that

$$R^{\mathrm{dt}}_{2^{-k-3}}(f_{n,k,U}) = \Omega\left(\frac{n^{1-\frac{1}{k}}}{(\log n)^{2-\frac{1}{k}}}\right).$$

By the error reduction formula (6.5.1), this settles (6.5.10) and (6.5.11). □

Theorem 6.44 settles Theorem 6.1 from the introduction. Corollary 6.2 now follows from (6.5.5) and Theorem 6.1 by taking $k = \lceil 1/\varepsilon \rceil + 1$ and $\gamma = 1/6$. Similarly, Corollary 6.3 follows from (6.5.5) and Theorem 6.1 by setting $\gamma = 1/6$ and taking $k = k(n)$ to be a sufficiently slow-growing function.

CHAPTER 7

Conclusion

In this dissertation, we studied communication with regard to the following questions: 1. how do we handle noise; 2. what is the communication complexity of certain computational tasks; and 3. what is the relationship between quantum and classical communication. We put these questions in concrete problems and were able to give very satisfying answers to them. We review these results one by one.

Communication against noise. In Chapter 4, we studied the BGMO corruption model of substitutions, insertions and deletions. We showed that there is an interactive coding scheme that uses a constant-size alphabet and achieves the optimal noise tolerance rate, at the expense of a constant-factor overhead in communication complexity compared to π.

There are two notable features of our coding scheme. First, we use the combinatorial structure of tree code. As discussed, the existence of a $(1-\delta)$-tree code was shown using probabilistic methods. Although we can always construct a $(1-\delta)$-tree code of depth n by exhaustive search, it is desirable to have a construction that is efficient. Progress in this direction is reported in [**23, 19, 61, 20, 60, 49**].

Open Problem 7.1. Is there an efficient way to construct a $(1-\delta)$-tree code for $\delta \in (0,1)$?

Second, to tolerate the maximum noise rate we used a large alphabet. What if we stick to a small-size alphabet? Assume the standard model, where every corruption is a substitution. When the size of the alphabet can be an arbitrary constant, every protocol has a coding scheme with a linear rate that can tolerate a corruption rate of $1/4 - \varepsilon$ for any $\varepsilon > 0$. For the binary alphabet, however, the optimum rate is only known to be between $\frac{5}{39} - \varepsilon$ and $\frac{1}{6}$ [**51, 52**].

Open Problem 7.2. What is the maximum corruption rate that can be tolerated using interactive coding with a constant code rate and a *binary* alphabet?

Sign-representations. In Chapter 5, we studied the sign-representations of $\mathbf{AC}^0$. We proved essentially optimal threshold degree and sign-rank lower bounds for $\mathbf{AC}^0$. These analytical lower bounds in turn imply our strong lower bounds of the communication complexity. Moreover, the techniques here are strong and have implications in other areas of theoretical computer science, like learning theory.

Our threshold degree result for $\mathbf{AC}^0$ took advantage of large depth circuits. If we turn to extremely shallow circuits, there are unsettled problems. Depth-1 circuits are just the AND and OR functions, which we understand very well. But analyzing the sign representation of depth-2 circuits is already very challenging. For example, an important open problem is the quantum query complexity of *triangle detection*—to decide if a graph contains a triangle. Contrary to the sophisticated quantum algorithms we have, there is no nontrivial lower bound. Establishing a tight approximate degree lower bound is currently the most promising approach to this fundamental problem. Formally, for any graph with vertices labeled by $1, 2, \ldots, n$, let e_{ij} indicates whether vertices i and j are connected. Then $\mathrm{TRI}_n : \{0,1\}^{n(n-1)/2} \to \{0,1\}$ is

$$\mathrm{TRI}_n(e) = \bigvee_{1 \le i < j < k \le n} e_{ij} \wedge e_{jk} \wedge e_{ik}.$$

Open Problem 7.3. Determine the approximate degree of the triangle detection problem.

Quantum query/communication complexity. In Chapter 6, we exhibited an optimal separation between quantum and classical query complexity with respect to partial functions. This separation then transfers to a near-optimal separation between quantum and classical communication complexity.

Now, in the query world, our understanding is more or less complete. We know that the quantum and randomized query complexity can be arbitrarily separated for partial functions. For total functions we have this cubic separation, and Aaronson et al. recently proved that the separation is at most quartic [**4**]. Closing this gap for total functions is an important open problem.

Open Problem 7.4. Prove or disprove that for any total function f,

$$R^{\mathrm{dt}}(f) \leq (Q^{\mathrm{dt}}(f))^3.$$

In communication world, things are more open, especially for total functions. It is a major open problem to decide if quantum protocols (even with prior entanglement) can be super-polynomially more efficient than randomized protocols for total functions. The reason why this problem is so challenging is that we lack tools for such a result. For example, the "lifting" technique will not help since quantum and randomized query complexity are polynomially related. Moreover, many lower bound techniques (e.g. approximate rank, discrepancy method) for randomized communication complexity also lower bound quantum communication complexity.

Open Problem 7.5. Is there a (family of) total function $f : \{0,1\}^n \to \{0,1\}$ such that $Q^*(f) = (R(f))^{o(1)}$?

Bibliography

[1] Scott Aaronson and Andris Ambainis. Quantum search of spatial regions. *Theory of Computing*, 1(1):47–79, 2005.

[2] Scott Aaronson and Andris Ambainis. Forrelation: A problem that optimally separates quantum from classical computing. *SIAM J. Comput.*, 47(3):982–1038, 2018.

[3] Scott Aaronson, Shalev Ben-David, and Robin Kothari. Separations in query complexity using cheat sheets. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing* (STOC), pages 863–876, 2016.

[4] Scott Aaronson, Shalev Ben-David, Robin Kothari, Shravas Rao, and Avishay Tal. Degree vs. approximate degree and quantum implications of Huang's sensitivity theorem. In *Proceedings of the Fifty-Third Annual ACM Symposium on Theory of Computing* (STOC), New York, NY, USA, 2021.

[5] Scott Aaronson and Yaoyun Shi. Quantum lower bounds for the collision and the element distinctness problems. *J. ACM*, 51(4):595–605, 2004.

[6] Shweta Agrawal, Ran Gelles, and Amit Sahai. Adaptive protocols for interactive communication. In *IEEE International Symposium on Information Theory, ISIT 2016, Barcelona, Spain, July 10-15, 2016*, pages 595–599, 2016.

[7] Noga Alon, Peter Frankl, and Vojtech Rödl. Geometrical realization of set systems and probabilistic communication complexity. In *Proceedings of the Twenty-Sixth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 277–280, 1985.

[8] Andris Ambainis. Polynomial degree and lower bounds in quantum complexity: Collision and element distinctness with small range. *Theory of Computing*, 1(1):37–46, 2005.

[9] Andris Ambainis, Andrew M. Childs, Ben Reichardt, Robert Špalek, and Shengyu Zhang. Any AND-OR formula of size N can be evaluated in time $N^{1/2+o(1)}$ on a quantum computer. *SIAM J. Comput.*, 39(6):2513–2530, 2010.

[10] James Aspnes, Richard Beigel, Merrick L. Furst, and Steven Rudich. The expressive power of voting polynomials. *Combinatorica*, 14(2):135–148, 1994.

[11] László Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *Proceedings of the Twenty-Seventh Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 337–347, 1986.

[12] László Babai, Noam Nisan, and Mario Szegedy. Multiparty protocols, pseudorandom generators for logspace, and time-space trade-offs. *J. Comput. Syst. Sci.*, 45(2):204–232, 1992.

[13] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and R. de Wolf. Quantum lower bounds by polynomials. *J. ACM*, 48(4):778–797, 2001.

[14] Paul Beame and Trinh Huynh. Multiparty communication complexity and threshold circuit size of AC^0. *SIAM J. Comput.*, 41(3):484–518, 2012.

[15] Paul Beame and Widad Machmouchi. The quantum query complexity of AC^0. *Quantum Information & Computation*, 12(7-8):670–676, 2012.

[16] Richard Beigel, Nick Reingold, and Daniel A. Spielman. PP is closed under intersection. *J. Comput. Syst. Sci.*, 50(2):191–202, 1995.

[17] Shai Ben-David, Nadav Eiron, and Hans Ulrich Simon. Limitations of learning via embeddings in Euclidean half spaces. *J. Mach. Learn. Res.*, 3:441–461, 2003.

[18] Ethan Bernstein and Umesh V. Vazirani. Quantum complexity theory. *SIAM J. Comput.*, 26(5):1411–1473, 1997.

[19] Zvika Brakerski and Yael Tauman Kalai. Efficient interactive coding against adversarial noise. In *Proceedings of the Fifty-Third Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 160–166, 2012.

[20] Zvika Brakerski, Yael Tauman Kalai, and Moni Naor. Fast interactive coding against adversarial noise. *J. ACM*, 61(6):35:1–35:30, 2014.

[21] Zvika Brakerski, Yael Tauman Kalai, and Raghuvansh R. Saxena. Deterministic and efficient interactive coding from hard-to-decode tree codes. In *Proceedings of the Sixty-First Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 446–457. IEEE, 2020.

[22] Gilles Brassard, Ashwin Nayak, Alain Tapp, Dave Touchette, and Falk Unger. Noisy interactive quantum communication. In *Proceedings of the Fifty-Fifth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 296–305, 2014.

[23] Mark Braverman. Towards deterministic tree code constructions. In *Innovations in Theoretical Computer Science 2012, Cambridge, MA, USA, January 8-10, 2012*, pages 161–167, 2012.

[24] Mark Braverman and Klim Efremenko. List and unique coding for interactive communication in the presence of adversarial noise. In *Proceedings of the Fifty-Fifth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 236–245, 2014.

[25] Mark Braverman, Ran Gelles, Jieming Mao, and Rafail Ostrovsky. Coding for interactive communication correcting insertions and deletions. *IEEE Trans. Information Theory*, 63(10):6256–6270, 2017.

[26] Mark Braverman and Anup Rao. Toward coding for maximum errors in interactive communication. *IEEE Trans. Information Theory*, 60(11):7248–7255, 2014.

[27] Sergey Bravyi, David Gosset, and Daniel Grier. Classical algorithms for forrelation. Available at http://arxiv.org/abs/2102.06963, 2021.

[28] Harry Buhrman, Richard Cleve, and Avi Wigderson. Quantum vs. classical communication and computation. In *Proceedings of the Thirtieth Annual ACM Symposium on Theory of Computing* (STOC), pages 63–68, 1998.

[29] Harry Buhrman and R. de Wolf. Complexity measures and decision tree complexity: A survey. *Theor. Comput. Sci.*, 288(1):21–43, 2002.

[30] Harry Buhrman, Lance Fortnow, Ilan Newman, and Hein Röhrig. Quantum property testing. *SIAM J. Comput.*, 37(5):1387–1400, 2008.

[31] Harry Buhrman, Ilan Newman, Hein Röhrig, and R. de Wolf. Robust polynomials and quantum algorithms. *Theory Comput. Syst.*, 40(4):379–395, 2007.

[32] Harry Buhrman, Nikolai K. Vereshchagin, and R. de Wolf. On computation and communication with small bias. In *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity* (CCC), pages 24–32, 2007.

[33] Mark Bun, Robin Kothari, and Justin Thaler. The polynomial method strikes back: Tight quantum query bounds via dual polynomials. ECCC Report TR17-169, 2017.

[34] Mark Bun and Justin Thaler. Dual lower bounds for approximate degree and Markov–Bernstein inequalities. *Inf. Comput.*, 243:2–25, 2015.

[35] Mark Bun and Justin Thaler. Hardness amplification and the approximate degree of constant-depth circuits. In *Proceedings of the Forty-Second International Colloquium on Automata, Languages and Programming* (ICALP), pages 268–280, 2015.

[36] Mark Bun and Justin Thaler. Approximate degree and the complexity of depth three circuits. In *Electronic Colloquium on Computational Complexity (ECCC)*, 2016. Report TR16-121.

[37] Mark Bun and Justin Thaler. Improved bounds on the sign-rank of $\mathbf{AC}^0$. In *Proceedings of the Forty-Third International Colloquium on Automata, Languages and Programming* (ICALP), pages 37:1–37:14, 2016.

[38] Mark Bun and Justin Thaler. A nearly optimal lower bound on the approximate degree of AC^0. In *Proceedings of the Fifty-Eighth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 1–12, 2017.

[39] Mark Bun and Justin Thaler. The large-error approximate degree of $\mathbf{AC}^0$. In *Electronic Colloquium on Computational Complexity (ECCC)*, August 2018. Report TR18-143.

[40] Ashok K. Chandra, Merrick L. Furst, and Richard J. Lipton. Multi-party protocols. In *Proceedings of the Fifteenth Annual ACM Symposium on Theory of Computing* (STOC), pages 94–99, 1983.

[41] Arkadev Chattopadhyay. Discrepancy and the power of bottom fan-in in depth-three circuits. In *Proceedings of the Forty-Eighth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 449–458, 2007.

[42] Arkadev Chattopadhyay and Anil Ada. Multiparty communication complexity of disjointness. In *Electronic Colloquium on Computational Complexity (ECCC)*, January 2008. Report TR08-002.

[43] Arkadev Chattopadhyay, Yuval Filmus, Sajin Koroth, Or Meir, and Toniann Pitassi. Query-to-communication lifting for BPP using inner product. In *Proceedings of the Forty-Sixth International Colloquium on Automata, Languages and Programming* (ICALP), volume 132 of *LIPIcs*, pages 35:1–35:15, 2019.

[44] Arkadev Chattopadhyay, Michal Koucký, Bruno Loff, and Sagnik Mukhopadhyay. Composition and simulation theorems via pseudo-random properties. *Comput. Complex.*, 28(4):617–659, 2019.

[45] Eshan Chattopadhyay, Pooya Hatami, Kaave Hosseini, and Shachar Lovett. Pseudorandom generators from polarizing random walks. In *Proceedings of the Thirty-Third Annual IEEE Conference on Computational Complexity* (CCC), volume 102, pages 1:1–1:21, 2018.

[46] Eshan Chattopadhyay, Pooya Hatami, Omer Reingold, and Avishay Tal. Improved pseudorandomness for unordered branching programs through local monotonicity. In *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing* (STOC), pages 363–375, 2018.

[47] Benny Chor and Oded Goldreich. Unbiased bits from sources of weak randomness and probabilistic communication complexity. *SIAM J. Comput.*, 17(2):230–261, 1988.

[48] Amit Choudhury and Paramita Roy. A fairly accurate approximation to the area under normal curve. *Commun. Stat. Simul. Comput.*, 38(7):1485–1492, 2009.

[49] Gil Cohen, Bernhard Haeupler, and Leonard J Schulman. Explicit binary tree codes with polylogarithmic size alphabet. In *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing* (STOC), pages 535–544, 2018.

[50] David Deutsch and Richard Jozsa. Rapid solution of problems by quantum computation. *Proc. R. Soc. Lond. A*, 439:553–558, 1992.

[51] Klim Efremenko, Ran Gelles, and Bernhard Haeupler. Maximal noise in interactive communication over erasure channels and channels with feedback. *IEEE Trans. Information Theory*, 62(8):4575–4588, 2016.

[52] Klim Efremenko, Gillat Kol, and Raghuvansh R. Saxena. Binary interactive error resilience beyond $^1/_8$ (or why $(^1/_2)^3 > {^1/_8}$). In *Proceedings of the Sixty-First Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 470–481, 2020.

[53] Jürgen Forster. A linear lower bound on the unbounded error probabilistic communication complexity. *J. Comput. Syst. Sci.*, 65(4):612–625, 2002.

[54] Jürgen Forster, Matthias Krause, Satyanarayana V. Lokam, Rustam Mubarakzjanov, Niels Schmitt, and Hans-Ulrich Simon. Relations between communication complexity, linear arrangements, and computational complexity. In *Proc. of the 21st Conf. on Foundations of Software Technology and Theoretical Computer Science (FST TCS)*, pages 171–182, 2001.

[55] Matthew K. Franklin, Ran Gelles, Rafail Ostrovsky, and Leonard J. Schulman. Optimal coding for streaming authentication and interactive communication. *IEEE Trans. Information Theory*, 61(1):133–145, 2015.

[56] Dmitry Gavinsky. Entangled simultaneity versus classical interactivity in communication complexity. *IEEE Trans. Inf. Theory*, 66(7):4641–4651, 2020.

[57] Dmitry Gavinsky, Julia Kempe, Oded Regev, and R. de Wolf. Bounded-error quantum state identification and exponential separations in communication complexity. In *Proceedings of the*

Thirty-Eighth Annual ACM Symposium on Theory of Computing (STOC), pages 594–603, 2006.

[58] Ran Gelles. Coding for interactive communication: A survey. *Foundations and Trends in Theoretical Computer Science*, 13(1-2):1–157, 2017.

[59] Ran Gelles and Bernhard Haeupler. Capacity of interactive communication over erasure channels and channels with feedback. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 1296–1311, 2015.

[60] Ran Gelles, Bernhard Haeupler, Gillat Kol, Noga Ron-Zewi, and Avi Wigderson. Towards optimal deterministic coding for interactive communication. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1922–1936, 2016.

[61] Ran Gelles, Ankur Moitra, and Amit Sahai. Efficient coding for interactive communication. *IEEE Trans. Information Theory*, 60(3):1899–1913, 2014.

[62] Mohsen Ghaffari and Bernhard Haeupler. Optimal error rates for interactive coding II: efficiency and list decoding. In *Proceedings of the Fifty-Fifth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 394–403, 2014.

[63] Mohsen Ghaffari, Bernhard Haeupler, and Madhu Sudan. Optimal error rates for interactive coding I: adaptivity and other settings. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing* (STOC), pages 794–803, 2014.

[64] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. In *Proceedings of the Fifty-Sixth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 1077–1088, 2015.

[65] Mika Goos, Toniann Pitassi, and Thomas Watson. Query-to-communication lifting for bpp. *SIAM Journal on Computing*, 49(4):FOCS17–441, 2020.

[66] Parikshit Gopalan, Rocco A. Servedio, and Avi Wigderson. Degree and sensitivity: Tails of two distributions. In *Proceedings of the Thirty-First Annual IEEE Conference on Computational Complexity* (CCC), volume 50, pages 13:1–13:23, 2016.

[67] P. Gordan. Über die Auflösung linearer Gleichungen mit reellen Coefficienten. *Mathematische Annalen*, 6:23–28, 1873.

[68] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, 2nd edition, 1994.

[69] Lov K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing* (STOC), pages 212–219, 1996.

[70] Venkatesan Guruswami and Ray Li. Efficiently decodable insertion/deletion codes for high-noise and high-rate regimes. In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 620–624, 2016.

[71] Bernhard Haeupler. Interactive channel capacity revisited. In *Proceedings of the Fifty-Fifth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 226–235, 2014.

[72] Hao Huang. Induced subgraphs of hypercubes and a proof of the sensitivity conjecture. *Annals of Mathematics*, 190(3):949–955, 2019.

[73] Stasys Jukna. *Extremal Combinatorics with Applications in Computer Science*. Springer-Verlag Berlin Heidelberg, 2nd edition, 2011.

[74] Jørn Justesen. Class of constructive asymptotically good algebraic codes. *IEEE Trans. Information Theory*, 18(5):652–656, 1972.

[75] Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM J. Discrete Math.*, 5(4):545–557, 1992.

[76] Hartmut Klauck. Lower bounds for quantum communication complexity. In *Proceedings of the Forty-Second Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 288–297, 2001.

[77] Adam R. Klivans, Ryan O'Donnell, and Rocco A. Servedio. Learning intersections and thresholds of halfspaces. *J. Comput. Syst. Sci.*, 68(4):808–840, 2004.

[78] Adam R. Klivans and Rocco A. Servedio. Learning DNF in time $2^{\tilde{O}(n^{1/3})}$. *J. Comput. Syst. Sci.*, 68(2):303–318, 2004.

[79] Adam R. Klivans and Rocco A. Servedio. Toward attribute efficient learning of decision lists and parities. *J. Machine Learning Research*, 7:587–602, 2006.

[80] Gillat Kol and Ran Raz. Interactive channel capacity. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing* (STOC), pages 715–724, 2013.

[81] Matthias Krause and Pavel Pudlák. On the computational power of depth-2 circuits with threshold and modulo gates. *Theor. Comput. Sci.*, 174(1–2):137–156, 1997.

[82] Matthias Krause and Pavel Pudlák. Computing Boolean functions by polynomials and threshold circuits. *Comput. Complex.*, 7(4):346–370, 1998.

[83] Ilan Kremer. Quantum communication. Master's thesis, Hebrew University, Computer Science Department, 1995.

[84] Eyal Kushilevitz and Noam Nisan. *Communication complexity.* Cambridge University Press, 1997.

[85] Troy Lee. A note on the sign degree of formulas, 2009. Available at `http://arxiv.org/abs/0909.4607`.

[86] Troy Lee and Adi Shraibman. Lower bounds in communication complexity. *Foundations and Trends in Theoretical Computer Science*, 3(4):263–399, 2009.

[87] Vladimir I. Levenshtein. Binary codes capable of correcting deletions, insertions, and reversals. *Soviet Physics Doklady*, 10(8):707–710, 1966.

[88] Nati Linial, Shahar Mendelson, Gideon Schechtman, and Adi Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439–463, 2007.

[89] Nati Linial and Adi Shraibman. Lower bounds in communication complexity based on factorization norms. *Random Struct. Algorithms*, 34(3):368–394, 2009.

[90] Marvin L. Minsky and Seymour A. Papert. *Perceptrons: An Introduction to Computational Geometry.* MIT Press, Cambridge, Mass., 1969.

[91] Ilan Newman. Private vs. common random bits in communication complexity. *Inf. Process. Lett.*, 39(2):67–71, 1991.

[92] Noam Nisan. CREW PRAMs and decision trees. *SIAM J. Comput.*, 20(6):999–1007, 1991.

[93] Noam Nisan and Mario Szegedy. On the degree of Boolean functions as real polynomials. *Computational Complexity*, 4:301–313, 1994.

[94] Ryan O'Donnell. *Analysis of Boolean Functions.* Cambridge University Press, 2014.

[95] Ryan O'Donnell and Rocco A. Servedio. Learning monotone decision trees in polynomial time. *SIAM J. Comput.*, 37(3):827–844, 2007.

[96] Ryan O'Donnell and Rocco A. Servedio. Extremal properties of polynomial threshold functions. *J. Comput. Syst. Sci.*, 74(3):298–312, 2008.

[97] Ryan O'Donnell and Rocco A. Servedio. New degree bounds for polynomial threshold functions. *Combinatorica*, 30(3):327–358, 2010.

[98] Rafail Ostrovsky, Yuval Rabani, and Leonard J. Schulman. Error-correcting codes for automatic control. *IEEE Trans. Information Theory*, 55(7):2931–2941, 2009.

[99] Ramamohan Paturi. On the degree of polynomials that approximate symmetric Boolean functions. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Theory of Computing* (STOC), pages 468–474, 1992.

[100] Ramamohan Paturi and Michael E. Saks. Approximating threshold circuits by rational functions. *Inf. Comput.*, 112(2):257–272, 1994.

[101] Ramamohan Paturi and Janos Simon. Probabilistic communication complexity. *J. Comput. Syst. Sci.*, 33(1):106–123, 1986.

[102] Anup Rao and Amir Yehudayoff. *Communication Complexity: and Applications*. Cambridge University Press, 2020.

[103] Ran Raz. Exponential separation of quantum and classical communication complexity. In *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing* (STOC), pages 358–367, 1999.

[104] Ran Raz and Avishay Tal. Oracle separation of BQP and PH. In *Proceedings of the Fifty-First Annual ACM Symposium on Theory of Computing* (STOC), pages 13–23, 2019.

[105] Alexander A. Razborov. Quantum communication complexity of symmetric predicates. *Izvestiya: Mathematics*, 67(1):145–159, 2003.

[106] Alexander A. Razborov and Alexander A. Sherstov. The sign-rank of AC^0. *SIAM J. Comput.*, 39(5):1833–1855, 2010. Preliminary version in *Proceedings of the Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2008.

[107] Oded Regev and Bo'az Klartag. Quantum one-way communication can be exponentially stronger than classical communication. In *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (STOC), pages 31–40, 2011.

[108] Michael E. Saks. Slicing the hypercube. *Surveys in Combinatorics*, pages 211–255, 1993.

[109] Leonard J. Schulman. Communication on noisy channels: A coding theorem for computation. In *Proceedings of the Thirty-Third Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 724–733, 1992.

[110] Leonard J. Schulman. Deterministic coding for interactive communication. In *Proceedings of the Twenty-Fifth Annual ACM Symposium on Theory of Computing* (STOC), pages 747–756, 1993.

[111] Leonard J. Schulman. Coding for interactive communication. *IEEE Trans. Information Theory*, 42(6):1745–1756, 1996.

[112] Leonard J. Schulman and David Zuckerman. Asymptotically good codes correcting insertions, deletions, and transpositions. *IEEE Trans. Information Theory*, 45(7):2552–2557, 1999.

[113] Alexander A. Sherstov. Halfspace matrices. *Computational Complexity*, 17(2):149–178, 2008. Preliminary version in *Proceedings of the Twenty-Second Annual IEEE Conference on Computational Complexity* (CCC), 2007.

[114] Alexander A. Sherstov. Separating AC^0 from depth-2 majority circuits. *SIAM J. Comput.*, 38(6):2113–2129, 2009. Preliminary version in *Proceedings of the Thirty-Ninth Annual ACM Symposium on Theory of Computing* (STOC), 2007.

[115] Alexander A. Sherstov. Communication complexity under product and nonproduct distributions. *Computational Complexity*, 19(1):135–150, 2010. Preliminary version in *Proceedings of the Twenty-Third Annual IEEE Conference on Computational Complexity* (CCC), 2008.

[116] Alexander A. Sherstov. The pattern matrix method. *SIAM J. Comput.*, 40(6):1969–2000, 2011. Preliminary version in *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing* (STOC), 2008.

[117] Alexander A. Sherstov. The unbounded-error communication complexity of symmetric functions. *Combinatorica*, 31(5):583–614, 2011. Preliminary version in *Proceedings of the Forty-Ninth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2008.

[118] Alexander A. Sherstov. Strong direct product theorems for quantum communication and query complexity. *SIAM J. Comput.*, 41(5):1122–1165, 2012. Preliminary version in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing* (STOC), 2011.

[119] Alexander A. Sherstov. The intersection of two halfspaces has high threshold degree. *SIAM J. Comput.*, 42(6):2329–2374, 2013. Preliminary version in *Proceedings of the Fiftieth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2009.

[120] Alexander A. Sherstov. Making polynomials robust to noise. *Theory of Computing*, 9:593–615, 2013. Preliminary version in *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing* (STOC), 2012.

[121] Alexander A. Sherstov. Optimal bounds for sign-representing the intersection of two halfspaces by polynomials. *Combinatorica*, 33(1):73–96, 2013. Preliminary version in *Proceedings of the Forty-Second Annual ACM Symposium on Theory of Computing* (STOC), 2010.

[122] Alexander A. Sherstov. Breaking the Minsky–Papert barrier for constant-depth circuits. In *Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing* (STOC), pages 223–232, 2014. Full version available as ECCC Report TR14-009, January 2014.

[123] Alexander A. Sherstov. Communication lower bounds using directional derivatives. *J. ACM*, 61(6):1–71, 2014. Preliminary version in *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing* (STOC), 2013.

[124] Alexander A. Sherstov. The power of asymmetry in constant-depth circuits. In *Proceedings of the Fifty-Sixth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 431–450, 2015.

[125] Alexander A. Sherstov. The multiparty communication complexity of set disjointness. *SIAM J. Comput.*, 45(4):1450–1489, 2016. Preliminary version in *Proceedings of the Forty-Fourth Annual ACM Symposium on Theory of Computing* (STOC), 2009.

[126] Alexander A. Sherstov. Algorithmic polynomials. In *Proceedings of the Fiftieth Annual ACM Symposium on Theory of Computing* (STOC), pages 311–324, 2018.

[127] Alexander A. Sherstov, Andrey A. Storozhenko, and Pei Wu. An optimal separation of randomized and quantum query complexity. In *Proceedings of the Fifty-Third Annual ACM Symposium on Theory of Computing* (STOC), 2021.

[128] Alexander A. Sherstov and Pei Wu. Optimal interactive coding for insertions, deletions, and substitutions. In *Proceedings of the Fifty-Eighth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 240–251, 2017.

[129] Alexander A. Sherstov and Pei Wu. Near-optimal lower bounds on the threshold degree and sign-rank of AC^0. In *Proceedings of the Fifty-First Annual ACM Symposium on Theory of Computing* (STOC), pages 401–412, 2019.

[130] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.

[131] Daniel R. Simon. On the power of quantum computation. *SIAM J. Comput.*, 26(5):1474–1483, 1997.

[132] Kai-Yeung Siu, Vwani P. Roychowdhury, and Thomas Kailath. Rational approximation techniques for analysis of neural networks. *IEEE Transactions on Information Theory*, 40(2):455–466, 1994.

[133] Thomas Steinke, Salil P. Vadhan, and Andrew Wan. Pseudorandomness and Fourier-growth bounds for width-3 branching programs. *Theory Comput.*, 13(1):1–50, 2017.

[134] Avishay Tal. Tight bounds on the fourier spectrum of AC0. In *Proceedings of the Thirty-Second Annual IEEE Conference on Computational Complexity* (CCC), volume 79, pages 15:1–15:31, 2017.

[135] Avishay Tal. Towards optimal separations between quantum and randomized query complexities. In *Proceedings of the Sixty-First Annual IEEE Symposium on Foundations of Computer Science* (FOCS), 2020.

[136] Justin Thaler. Lower bounds for the approximate degree of block-composed functions. In *Proceedings of the Forty-Third International Colloquium on Automata, Languages and Programming* (ICALP), pages 17:1–17:15, 2016.

[137] Robert Špalek. A dual polynomial for OR. Available at `http://arxiv.org/abs/0803.4516`, 2008.

[138] Ronald de Wolf. *Quantum Computing and Communication Complexity.* PhD thesis, University of Amsterdam, 2001.

[139] Andrew Chi-Chih Yao. Some complexity questions related to distributive computing. In *Proceedings of the Eleventh Annual ACM Symposium on Theory of Computing* (STOC), pages 209–213, 1979.

[140] Andrew Chi-Chih Yao. Quantum circuit complexity. In *Proceedings of the Thirty-Fourth Annual IEEE Symposium on Foundations of Computer Science* (FOCS), pages 352–361, 1993.